arXiv:1511.06849v1 [stat.AP] 21 Nov 2015

Practical survival analysis tools for heterogeneous cohorts and informative censoring

M. Rowley[a], H. Garmo[b], M. van Hemelrijck[b], W. Wulaningsih[b], B. Grundmark[c,d], B. Zethelius[e,d], N. Hammar[f,g], G. Walldius[h], M. Inoue[i], L. Holmberg[b] and A.C.C. Coolen[a]

November 24, 2015

Abstract

In heterogeneous cohorts and those where censoring by non-primary risks is informative many conventional survival analysis methods are not applicable; the proportional hazards assumption is usually violated at population level and the observed crude hazard rates are no longer estimators of what they would have been in the absence of other risks. In this paper, we develop a fully Bayesian survival analysis to determine the probabilistically optimal description of a heterogeneous cohort and we propose a novel means of recovering hazard rates and survival functions 'decontaminated' of the effects of any competing risks. Most competing risks studies implicitly assume that risk correlations are induced by cohort or disease heterogeneity that is not captured by covariates. We additionally assume that proportional hazards hold at the level of individuals, for all risks, leading to a generic statistical description that allows us to decontaminate the effects of informative censoring, and from which Cox regression, frailty and random effects models, and latent class models can all be recovered as special cases. Synthetic data confirm that our approach can map a cohort's substructure, and remove heterogeneity-induced false protectivity and false aetiology effects. Application to survival data from the ULSAM cohort leads to plausible alternative explanations for previous counter-intuitive inferences to prostate cancer. The importance of managing cardiovascular disease as a comorbidity in women diagnosed with breast cancer is suggested on application to survival data from the AMORIS study.

Keywords: survival analysis; heterogeneity; informative censoring; competing risks

[a] Institute for Mathematical and Molecular Biomedicine, King's College London, London, U.K.
[b] Cancer Epidemiology Group, King's College London, School of Medicine, Guy's Hospital, London, U.K.
[c] Department of Surgical Sciences, Uppsala University, Uppsala, Sweden.
[d] Medical Products Agency, Uppsala, Sweden.
[e] Department of Public Health and Caring Sciences/Geriatrics, Uppsala University, Sweden.
[f] Department of Epidemiology, Institute of Environmental Medicine, Karolinska Institutet, Sweden.
[g] AstraZeneca Sverige, Södertalje, Sweden.
[h] Department of Cardiovascular Epidemiology, Institute of Environmental Medicine, Karolinska Institutet, Sweden.
[i] Department of Electrical Engineering and Bioscience, Waseda University, Tokyo, Japan.

Mark Rowley, Institute for Mathematical and Molecular Biomedicine, King's College London, Hodgkin Building, London SE1 1UL, U.K. E-mail: mark.rowley@kcl.ac.uk

Contract/grant sponsor: Prostate Cancer UK, European Union FP-7 Programme (IMAGINT), and the Ana Leaf Foundation

1 Introduction

The analysis of survival data is often complicated by cohort heterogeneity and informative censoring arising from competing risks. In this paper, we extend the conventional approaches to modelling a cohort and present a survival analysis which determines probabilistically the most likely characterisation of a cohort and can provide cause-specific hazard rates and survival curves 'decontaminated' of the effects of informative censoring.

A cohort is subject to informative censoring if the event-times of the primary and non-primary risks are not statistically independent; a cohort for which this is true is described as being subject to 'competing risks'. Unfortunately, one cannot infer presence or absence of risk correlations from survival data alone [1, 2], and in many cases the independence assumption which underpins many survival analysis approaches is expected to be incorrect. Residual heterogeneity, which is not visible in the covariates, is often a fingerprint of such risk event-time correlations in cohorts having competing risks. The inference of risk characteristics from survival data in the hypothetical situation where all other risks were disabled requires that the competing risk problem be addressed [3], that is, it is necessary to *somehow* handle contamination by informative censoring. In discussing informative censoring in a cohort, we shall refer to the *crude* cause-specific hazard rates and survival functions, in which the influence of competing risks is present, and we define their *decontaminated* analogues as the cause-specific decontaminated hazard rates and survival functions, isolated from the effects of all other risks.

The effective determination of heterogeneity can lead to a more accurate understanding of the characteristics of a cohort. In clinical epidemiology, diagnosing and disentangling cohort heterogeneity is crucial in many studies where differences in "case-mix" are considered to be problematic, e.g. in comparing treatment outcomes from different institutions. Unaccounted for risk correlations can lead to incorrect inferences [4, 5, 6, 7, 8] and the importance of having reliable epidemiological tools for isolating statistical features even for interrelated comorbid diseases is increasingly recognised [9]. An understanding of how a modifiable covariate (e.g. taking a particular medication, consuming a particular food type, smoking, exercising, etc.) having different influences for different groups within a cohort could prove advantageous for effective monitoring and management of population health. In populations where longevity is increasing, once relatively uncommon conditions are becoming more prevalent with the development of more effective therapies to treat competing conditions. Knowledge of the expected survival against such currently uncommon diseases, decontaminated of the influence of other risks, may provide a useful means of allocating research and public health resources efficiently.

Simple survival analysis methods, such as Kaplan-Meier estimators [10] and Cox regression [11], are unable to correctly capture the characteristics of a cohort in the presence of heterogeneity as their successful use requires proportional hazards across the cohort as a whole. The uncritical use of these techniques can be quite misleading when the proportional hazards assumption is not satisfied, as demonstrated in Figure 1 (a) for a simulated cohort modelling false protectivity effects. However, they can offer clues as to the presence of heterogeneity and informative censoring in a cohort. Inspection of the covariate-conditioned risk-specific Kaplan-Meier estimators can confirm violation of the proportional hazards assumption, as is clearly evident for the cohort in Figure 1 (b). Differences between the cumulative incidence [12] and the complement of the corresponding risk-specific Kaplan-Meier estimator are indicative of competing risks in a cohort e.g. [13], as shown in Figure 1 (c).

Many authors have tried to model residual cohort heterogeneity, usually starting from Cox-type cause-specific hazard rates, but with additional individualised risk multipliers. If the multipliers do not depend on the covariates we speak of 'frailty models', e.g. [14, 15, 16, 17, 18], and regard them as representing the impact of unobserved covariates, see e.g. [19]. If they depend on the covariates we would speak of 'random effects models', e.g. [20, 5, 21, 22, 23]. If the distribution of frailty factors takes the form of discrete clusters (latent classes, [24]), we obtain the latent class models; see e.g. [25] or [26] (which combines frailty and random effects with covariate-dependent class allocation as in [27]). Further variations include time-dependent frailty factors, and models in which the latent class of each individual is known. Most frailty and random effects studies, however, quantify only the hazard rate of the primary risk. They thereby capture some consequences of cohort heterogeneity, but without modelling also the non-primary risks it is fundamentally impossible to deal with the competing risk problem.

The approach of [28] focuses on parametrising the covariate-conditioned cumulative incidence function of the primary risk. It is conceptually similar to [11]; both model the primary risk profile in the presence of all risks. Cumulative incidence functions appear more intuitive than hazard rates; they are directly measurable, and incorporate also the impact of non-primary risks. However, expressing the data likelihood in terms of cumulative incidence functions is more cumbersome than in terms of hazard rates. But while [28] quantify risks that compete, they do not address the competing risk problem. Further developments involve e.g. alternative parametrisations [29, 30], application to the cumulative incidence of non-primary risks [31], and the inclusion of frailty factors [32].

Another community of authors have focused further on identifying which mathematical constraints or conditions need to be imposed on multi-risk survival analysis models in order to circumvent the identifiability problem, and

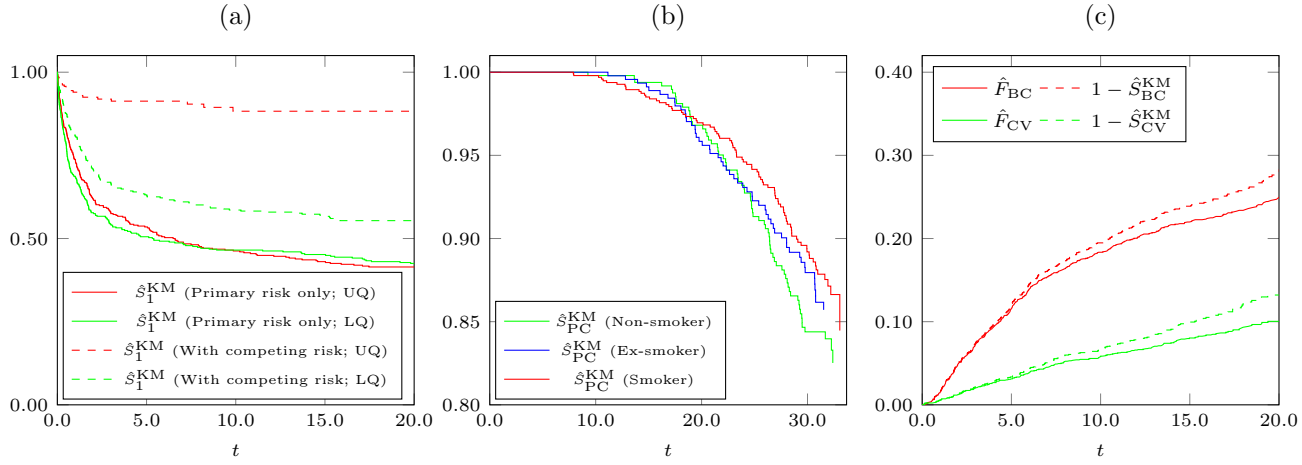

Figure 1: *Effects and signatures of heterogeneity and informative censoring:* In (a) an illustration of the dangers of using covariate-conditioned risk-specific Kaplan-Meier (KM) estimators, $\hat{S}_r^{\mathrm{KM}}$, in the presence of competing risks. The estimators for the lower and upper quartiles (LQ and UQ respectively) of a covariate associated with the primary risk are shown for two simulated cohorts, one in which there is no informative censoring and another which is subject to a competing risk, with both cohorts sharing *identical* primary risk characteristics. These estimators offer an extremely misleading indication of survival against the primary risk in the presence of informative censoring (i.e. a competing risk), as can be seen by the differences between the estimators in the absence (solid lines) and presence (dashed lines) of informative censoring. The application of our practical tools to provide *decontaminated* survival estimators, free from the effects of the competing risk, to the synthetic cohort data shown here is detailed in Section 3. In (b) the use of the covariate-conditioned risk-specific Kaplan-Meier estimator to detect the violation of proportional hazards assumption and infer possible heterogeneity in a cohort. The Kaplan-Meier estimator for prostate cancer (r =PC) risk, conditioned on an individuals smoking status, is suggestive of heterogeneity in the ULSAM cohort; the estimator would suggest that the survival of smokers is greater than that of ex- and non-smokers (an alternative view is suggested from our analysis; Section 4). In (c) the empirically measured cumulative incidence, $\hat{F}_r(t)$, and the complement of the Kaplan-Meier estimator, $\hat{S}_r^{\mathrm{KM}}$, for breast cancer (r =BC) and cardiovascular disease (r =CV) death for women diagnosed with breast cancer from the AMORIS population; the difference between the cumulative incidence and the complement of the corresponding Kaplan-Meier estimator indicating the presence of competing risks in the cohort.

infer the joint event time distribution unambiguously from survival data. Examples involving survival data with covariates are [33], and [34]. However, also these studies do not take the step towards decontamination tools or facilitate interpretation.

In our analysis, we build upon the frailty, random effects, and latent class approaches, to develop a generic model from which we estimate the relative frailty, covariate association(s), and base hazard rate for *each* latent class and for *all* risks. Additionally, by assuming that informative censoring at the cohort level is a consequence of residual (disease- or patient-) heterogeneity that is not captured by covariates in populations where only at the level of individuals are the different risks independent, we are able to derive exact formulae for decontaminated hazard rates and cause-specific survival functions. Our assumption of such heterogeneity-induced informative censoring is much weaker than assuming risk independence, yet it still imposes sufficient constraints to decontaminate from the effects of informative censoring. Furthermore, our analysis offers improvement over existing approaches in several distinct ways: In choosing to model *all* risks *simultaneously* more information can be extracted from cohort data in which risk correlations are present; in such circumstances, the application of an analysis approach which models the primary risk only is akin to choosing to not fully utilise the information available from the cohort data. Our analysis also introduces a fully Bayesian model selection for the determination of the optimal characterisation of a cohort.

In Section 2 we classify the distinct levels of "risk complexity" in a cohort from the competing risk perspective and define precisely what we mean by heterogeneity-induced informative censoring before presenting the mathematical development behind our approach to survival analysis.

In Section 3 the effectiveness of our analysis to characterise heterogeneous cohorts is demonstrated; the results obtained on application of our analysis to synthetic survival data, simulating a variety of heterogeneous cohorts and informative censoring, are presented and compared with those obtained from Kaplan-Meier estimators and Cox regression.

In Sections 4 and 5 our analysis is applied to real survival data related to prostate and breast cancer, from the ULSAM longitudinal cohort [35, 36] and AMORIS study [37, 38], respectively. The application to the ULSAM data leads to appealing and transparent new explanations for previously counter-intuitive inferences. Age-related survival differences between women diagnosed with breast cancer were found on application to data from the AMORIS cohort. In Section 6 we summarise our findings.

2 Heterogeneity-induced informative censoring

In this section we introduce our survival analysis based on the assumption that risk event time correlations, if present, are caused by residual cohort heterogeneity. This assumption is much weaker than that of assuming risk independence, but it does allow us to overcome informative censoring and leads to an intuitive and transparent parametrisation of hazard rates, and makes quantities *decontaminated* of its effects accessible.

The relationships between the event probabilities in a cohort and those of its members are critical to any analysis that seeks to address the issue of heterogeneity and are formalised in Section 2.1. In Section 2.2 the extent to which different risk (in)dependence assumptions limit survival analysis is explored; existing survival analysis methods and our heterogeneity-induced competing risk approach are compared in terms of the "risk complexity" of cohorts to which they can be applied. Decontaminated cause-specific survival functions and hazard rates are presented in Section 2.3, the difference between the decontaminated quantities and their respective crude counterparts providing a means of quantifying the severity of informative censoring in a cohort. The development of our latent class approach to modelling heterogeneous cohorts is outlined in Section 2.4 (with further details and identities given in Appendix A) and details of our Bayesian approach to the determination of the optimal characterisation of a cohort are given in Section 2.5. In Section 2.6 the various outputs of our analysis and their benefits are discussed.

2.1 Connection between cohort level and individual level descriptions

The standard mathematical relations of survival analysis are derived directly from the joint event time distribution and hold irrespective of whether we have a large or small cohort, or even a single individual. Below we formalise a number of quantities of interest, both for the individual and for the cohort as a whole, and the relationships between them; the index i is used to denote those personalised quantities specific to individual i, with no such index being present for cohort-specific quantities.

We imagine a cohort of N individuals who are each subject to R true risks, labelled by $r = 1 \ldots R$, and an end-of-trial censoring event denoted by $r = 0$[1]. The joint event time distribution $\mathcal{P}(t_0, \ldots, t_R)$ describes, for the cohort as a whole, the probability of the event times $(t_0, \ldots, t_R)$, where $t_r \geq 0$ is the time at which risk r triggers an event[2], with the personalised event time distribution being denoted by $\mathcal{P}_i(t_0, \ldots, t_R)$. The cohort level risk event time distribution is a direct average of the personalised event risk event time distributions, such that $\mathcal{P}(t_0, \ldots, t_R) = N^{-1} \sum_i \mathcal{P}_i(t_0, \ldots, t_R)$. It is an estimator of the joint event time probability for the greater population if the cohort is representative. Throughout this paper we shall assume that N is sufficiently large to ensure that the differences between $\mathcal{P}(t_0, \ldots, t_R)$ and the population's event time distribution are immaterial and that inferences made about the cohort are also true for the population.

In the interest of readability it should be assumed that products and summations over the risks run over the end-of-trial censoring risk and the true risks, such that $r = 0, \ldots, R$, shall be written as $\prod_r$ and $\sum_r$ respectively; should a particular risk r' be excluded then the product shall be written as $\prod_{r \neq r'}$. Similarly, summations over the cohort shall run over all members, labelled by $i = 1, \ldots, N$, are written as $\sum_i$.

The crude cause-specific hazard rate, that is the probability per unit time that given failures occur at time t if until then none of the possible events has yet occurred, follows from the joint event time distributions. The personalised cause-specific hazard rate, $h_r^i(t)$, is given by,

$$h_r^i(t) \;=\; \frac{1}{S_i(t)} \int_0^\infty \ldots \int_0^\infty \mathrm{d}t_0 \ldots \mathrm{d}t_R \; \mathcal{P}_i(t_0, \ldots, t_R) \delta(t - t_r) \prod_{r' \neq r} \theta(t_{r'} - t), \tag{1}$$

where the delta-distribution $\delta(x)$ is defined by the identity $\int_{-\infty}^{\infty} \mathrm{d}x \; \delta(x) f(x) = f(0)$, the step function is defined as $\theta(x > 0) = 1$ and $\theta(x < 0) = 0$, and the personalised survival function $S_i(t)$ is given by $S_i(t) = \mathrm{e}^{-\sum_r \int_0^t \mathrm{d}s \; h_r^i(s)}$. The personalised probability density, $P_i(t, r)$, to find the earliest event occurring at time t and corresponding to risk r is given by $P_i(t, r) = h_r^i(t) \mathrm{e}^{-\sum_{r'} \int_0^t \mathrm{d}s \; h_{r'}^i(s)}$.

At the cohort-level, analogous relations define the cause-specific hazard rate, $h_r(t)$, survival function, $S(t)$, and time-risk event probability, $P(t, r)$, for the cohort as a whole; these relations are of identical form to the relations above though obviously cohort-level quantities replace individual-specific quantities. Typically, however, it is the characterisation of the cohort *given* the survival analysis data for the cohort that is of interest; to achieve this we

[1] For the structure of the theory there is no difference between censoring due to alternative risks and censoring due to trial termination.

[2] This starting point is not fully general. It assumes that all risks will ultimately lead to failure. One can include events with a finite chance of not happening at *any* time, by adding for each risk r a binary variable τ_r to indicate whether or not the calamity button is pressed at time t_r.

require the covariate-conditioned cause-specific hazard rates and survival functions. The covariate-conditioned risk event time probability distribution $\mathcal{P}(t_0,\ldots,t_R|\mathbf{z})$ describes the risk event time statistics of the sub-cohort of those individuals i that have covariate vector $\mathbf{z}_i=\mathbf{z}$. At the cohort level, the covariate-conditioned cause-specific hazard rate is given by,

$$h_r(t|\mathbf{z}) = \frac{1}{S(t|\mathbf{z})}\int_0^\infty \ldots \int_0^\infty \mathrm{d}t_0 \ldots \mathrm{d}t_R\ \mathcal{P}(t_0,\ldots,t_R|\mathbf{z})\delta(t-t_r)\prod_{r'\neq r}\theta(t_{r'}-t), \tag{2}$$

where the covariate-conditioned survival function is given by $S(t|\mathbf{z}) = \mathrm{e}^{-\sum_{r'=0}^{R}\int_0^t \mathrm{d}s\ h_{r'}(s|\mathbf{z})}$. The covariate-conditioned risk event time probability density is then given by $P(t,r|\mathbf{z}) = h_r(t|\mathbf{z})\mathrm{e}^{-\sum_{r'=0}^{R}\int_0^t \mathrm{d}s\ h_{r'}(s|\mathbf{z})}$.

The connection between the cohort-level and personalised risk event time distributions is simple; accounting for covariate-conditioning, the cohort level risk event time distribution is the average of the personalised event time distributions of those individials with $\mathbf{z}_i=\mathbf{z}$, such that $\mathcal{P}(t_0,\ldots,t_R|\mathbf{z}) = \sum_{i,\ \mathbf{z}_i=\mathbf{z}}\mathcal{P}_i(t_0,\ldots,t_R)/\sum_{i,\ \mathbf{z}_i=\mathbf{z}}1$.

The relationship between quantities at the cohort and individual level is simple for those which depend linearly on the risk event time distribution. The cohort-level survival function is simply the average over the cohort of the individual survival functions, $S(t|\mathbf{z}) = \sum_{i,\ \mathbf{z}_i=\mathbf{z}} S_i(t)/\sum_{i,\ \mathbf{z}_i=\mathbf{z}}1$ and the risk event time probability density is $P(t,r|\mathbf{z}) = \sum_{i,\ \mathbf{z}_i=\mathbf{z}} P_i(t,r)/\sum_{i,\ \mathbf{z}_i=\mathbf{z}}1$. In contrast, quantities which depend on the risk event time distribution in a more complicated way, such as the crude cause-specific hazard rates, and cohort-level quantities are *not* direct averages over their individual level counterparts. The cohort level cause-specific hazard rates, for instance, are given by (see appendix A.1 for details),

$$h_r(t|\mathbf{z}) \quad = \quad \frac{\sum_{i,\mathbf{z}_i=\mathbf{z}} h_r^i(t)\mathrm{e}^{-\sum_{r'=0}^{R}\int_0^t \mathrm{d}s\ h_{r'}^i(s)}}{\sum_{i,\mathbf{z}_i=\mathbf{z}}\mathrm{e}^{-\sum_{r'=0}^{R}\int_0^t \mathrm{d}s\ h_{r'}^i(s)}}. \tag{3}$$

The cause-specific cumulative incidence function, $F_r(t)$, describes the probability that event r has been *observed* at any time prior to time t, and is given by,

$$F_r(t) \quad = \quad \int_0^t \mathrm{d}t'\ S(t')h_r(t'). \tag{4}$$

Although $F_r(t)$ refers to risk r specifically, it can be heavily influenced by other risks. If it is small, this may be because event r is intrinsically unlikely, or because it tends to be preceded by events $r'\neq r$. One cannot tell.

2.2 Risk complexity in heterogeneous cohorts

In this section we explore the the consequences of event time risk correlations and cohort-level heterogeneity. One always allows cohorts to be heterogeneous in terms of covariates; we refer here to heterogeneity in the relation between covariates and risks.

A homogeneous cohort is one in which the relationship between an individual's covariates and risk is uniform throughout the cohort. In a homogeneous cohort the personalised event time distribution $\mathcal{P}_i(t_0,\ldots,t_R)$ can depend on i only via $\mathbf{z}_i$ as such there exists a function $\mathcal{Q}(t_0,\ldots,t_R,\mathbf{z})$ such that $\mathcal{P}_i(t_0,\ldots,t_R) = \mathcal{Q}(t_0,\ldots,t_R,\mathbf{z}_i)$ is true for all individuals in the cohort. In such a cohort there can be no informative censoring and the crude and true cause-specific hazard rates and survival functions are identical.

In heterogeneous cohorts individuals have further relevant features which are not captured by their covariates. A consequence of these further features is that for such cohorts there no longer exists a function which, for all individuals in the cohort, maps between an individual's covariates and their personalised event time distribution. These additional features impact upon their risks. Here one will observe a gradual 'filtering': high-risk individuals will drop out early, causing time dependencies at cohort level that have no counterpart at individual level. For instance, even if all individuals have stationary hazard rates, one would according to (3) still find time dependent crude cohort level hazard rates. Here, having uncorrelated individual level risks no longer implies having uncorrelated covariate-conditioned cohort level risks. It is quite possible to have $\mathcal{P}_i(t_0,\ldots,t_R) = \prod_r \mathcal{P}_i(t_r)$, but still $\mathcal{P}(t_0,\ldots,t_R|\mathbf{z}) \neq \prod_r \mathcal{P}(t_r|\mathbf{z})$.

Risk event time correlations in a cohort can be generated at different levels. There is a natural hierarchy of cohorts in terms of risk complexity, as summarised in Table 1, with implications for the applicability of the different survival analysis methods.

Table 1: Risk complexity to capture cohort heterogeneity, risk event-time correlations, and competing risks

Risk complexity		Individual	Cohort
I	Homogenous cohort No competing risks at individual level No competing risks at cohort level	$\mathcal{P}_i(t_0,\ldots,t_R) = \prod_r \mathcal{P}(t_r\|\mathbf{z}_i)$	$\mathcal{P}(t_0,\ldots,t_R\|\mathbf{z}) = \prod_r \mathcal{P}(t_r\|\mathbf{z})$
II	Heterogeneous cohort No competing risks at individual level No competing risks at cohort level	$\mathcal{P}_i(t_0,\ldots,t_R) = \prod_r \mathcal{P}_i(t_r)$	$\mathcal{P}(t_0,\ldots,t_R\|\mathbf{z}) = \prod_r \mathcal{P}(t_r\|\mathbf{z})$
III	Heterogeneity-induced competing risks No competing risks at individual level Cohort level competing risks	$\mathcal{P}_i(t_0,\ldots,t_R) = \prod_r \mathcal{P}_i(t_r)$	$\mathcal{P}(t_0,\ldots,t_R\|\mathbf{z}) \neq \prod_r \mathcal{P}(t_r\|\mathbf{z})$
IV	Heterogeneous cohort Individual level competing risks Cohort level competing risks	$\mathcal{P}_i(t_0,\ldots,t_R) \neq \prod_r \mathcal{P}_i(t_r)$	$\mathcal{P}(t_0,\ldots,t_R\|\mathbf{z}) \neq \prod_r \mathcal{P}(t_r\|\mathbf{z})$

Assuming statistically independent risk event times at cohort level underlies both Kaplan-Meier estimators and Cox regression, and is seen only in those cohorts with risk complexity levels I and II. At level II there is still no competing risk problem, but heterogeneity may demand parametrisations of crude cohort level primary hazard rates that could be more complex than those of Cox, which is the rationale behind frailty and random effects models, and the latent class models of [26]. All these approaches still only model the primary risk, and therefore cannot handle cohorts beyond level II.

In this paper we shall focus on developing survival analysis tools to investigate cohorts with risk complexity level III, in which the event times of all risks are assumed to be statistically independent for each individual but residual heterogeneity leads to risk correlations at cohort level. In such cohorts the correlations between cohort level event times have their origin strictly in *correlations between disease susceptibilities and covariate associations of individuals*, for example, someone with a high hazard rate for a disease A may also be likely to have a high hazard rate for event B, for reasons not explained by the covariates. At this level of risk complexity competing risks phenomena shall be observed as the risk correlations at cohort level will cause informative censoring.

Cohorts of risk complexity level IV are the most complex and difficult to model, with the event times of different risks assumed to be correlated at both individual and cohort level, and shall not be explored in this paper.

2.3 Separating direct from indirect associations and quantifying informative censoring

The assumption of heterogeneity-induced competing risks, being that the event times of all risks are statistically independent for each individual in the cohort but not for the cohort as a whole (risk complexity level III in our classification above), allows us to investigate *analytically* the effects of informative censoring. This additional assumption addresses the issue of identifiabilty [1] and, with the risk event time marginal distributions now accessible, it becomes possible to develop expressions not only for the crude cause-specific hazard rates and survival function but also for their decontaminated counterparts.

The personalised cause-specific event time probability given the assumption of risk independence at the individual level is given by $\mathcal{P}_i(t_r) = h^i_r(t)\mathrm{e}^{-\int_0^t \mathrm{d}s\ h^i_r(s)}$. The covariate-conditioned cohort-level risk event time marginals are therefore given by $\mathcal{P}(t_r|\mathbf{z}) = \sum_{i,\mathbf{z}_i=\mathbf{z}} h^i_r(t)\mathrm{e}^{-\int_0^t \mathrm{d}s\ h^i_r(s)} / \sum_{i,\mathbf{z}_i=\mathbf{z}} 1$, and can be used to develop expressions for the decontaminated survival functions and hazard rates[3].

The crude cause-specific hazard rates which would have been found if all risks had been independent, $h_r(t|\mathbf{z})$, and their decontaminated counterparts, $\tilde{h}_r(t|\mathbf{z})$, are given by,

$$h_r(t|\mathbf{z}) = \frac{\sum_{i,\mathbf{z}_i=\mathbf{z}} h^i_r(t)\mathrm{e}^{-\sum_{r'=0}^R \int_0^t \mathrm{d}s\ h^i_{r'}(s)}}{\sum_{i,\mathbf{z}_i=\mathbf{z}} \mathrm{e}^{-\sum_{r'=0}^R \int_0^t \mathrm{d}s\ h^i_{r'}(s)}}, \qquad \tilde{h}_r(t|\mathbf{z}) = \frac{\sum_{i,\mathbf{z}_i=\mathbf{z}} h^i_r(t)\mathrm{e}^{-\int_0^t \mathrm{d}s\ h^i_r(s)}}{\sum_{i,\mathbf{z}_i=\mathbf{z}} \mathrm{e}^{-\int_0^t \mathrm{d}s\ h^i_r(s)}}. \tag{5}$$

The crude and decontaminated cause-specific hazard rates will generally have different values. In the decontaminated cause-specific hazard rate, $\tilde{h}_r(t|\mathbf{z})$, the probability that individual i survives until time t is given by $\exp[-\int_0^t \mathrm{d}s\ h^i_r(s)]$. The probability of survival until time t for individual i in the crude cause-specific hazard rate, $h_r(t|\mathbf{z})$, depends on *all* risks.

[3]The decontaminated survival functions and hazard rates follow from $\tilde{S}_r(t|\mathbf{z}) = \int_t^\infty \mathrm{d}t_r\ \mathcal{P}(t_r|\mathbf{z})$ and $\tilde{h}_r(t|\mathbf{z}) = -\frac{\mathrm{d}}{\mathrm{d}t}\log \tilde{S}_r(t|\mathbf{z})$.

Table 2: Personalised cause-specific hazard rates, for each of the R active risks, to capture latent cohort heterogeneity

$M = 1$	Heterogeneous frailties Homogeneous associations Homogeneous base hazard rates	$h_r^i(t) = \lambda_r(t)\mathrm{e}^{\beta_r^{\ell 0} + \sum_\mu \beta_r^\mu z_i^\mu}$
$M = 2$	Heterogeneous frailties Heterogeneous associations Homogeneous base hazard rates	$h_r^i(t) = \lambda_r(t)\mathrm{e}^{\beta_r^{\ell 0} + \sum_\mu \beta_r^{\ell\mu} z_i^\mu}$
$M = 3$	Heterogeneous frailties Heterogeneous associations Heterogeneous base hazard rates	$h_r^i(t) = \lambda_r^\ell(t)\mathrm{e}^{\beta_r^{\ell 0} + \sum_\mu \beta_r^{\ell\mu} z_i^\mu}$

The crude cause-specific survival function, $S_r(t|\mathbf{z})$, and its decontaminated counterpart, $\tilde{S}_r(t|\mathbf{z})$, are given by,

$$S_r(t|\mathbf{z}) = \mathrm{e}^{-\int_0^t \mathrm{d}s\ h_r(s|\mathbf{z})}, \qquad \tilde{S}_r(t|\mathbf{z}) = \frac{\sum_{i,\mathbf{z}_i=\mathbf{z}} \mathrm{e}^{-\int_0^t \mathrm{d}s\ h_r^i(s)}}{\sum_{i,\mathbf{z}_i=\mathbf{z}} 1}. \tag{6}$$

The decontaminated cause-specific survival function, $\tilde{S}_r(t|\mathbf{z})$, depends only on the risk r, whereas its crude counterpart depends on *all* risks through the crude hazard rate $h_r(t|\mathbf{z})$.

Hence the assumption that competing risks (if present) are induced by heterogeneity leads to relatively simple formulae for the decontaminated cause-specific quantities of interest and for the likelihood of observing individual survival data. What remains is to identify the *minimal* level of description required for evaluating these formulae, and to determine how the required information can be estimated from survival data.

2.4 Modelling the heterogeneous cohort

In this section we shall develop our latent class model. At the heart of our model are the personalised cause-specific hazard rates, $h_r^i(t)$, being of the Cox form and obeying the assumption of proportional hazards.

In constructing our latent class framework, we assume a heterogeneous cohort to be comprised of L sub-cohorts, or latent classes, labelled by $\ell = 1, \ldots, L$. Each individual from the cohort must belong to one of the latent classes; as such the classes are discrete, no class can contain a member from any other class, and together the classes contain all of the individuals in the cohort. Each class obeys the proportional hazards assumption although the cohort collectively need not.

The personalised cause-specific hazard rates shall depend on cause-specific and possibly class-specific frailty, association, and base hazard rate parameters. The frailty parameters, $\beta_r^{\ell 0}$, capture the impact of unobserved variables within class ℓ for risk r, i.e. effects that cannot be attributed to the included covariates directly, the association parameters, $\beta_r^{\ell\mu}$, quantify how strongly each of the $\mu = 1, \ldots, P$ covariates influence the personalised hazard rate, and the base hazard rate, $\lambda_r^\ell(t)$, describes the cause-specific personalised hazard rate for an individual from class ℓ having exactly *average* covariate values. Inevitably, modelling *all* risks and their correlations leads to models having more parameters than those which model *only* the primary risk. In an effort to avoid overfitting, three variants of the personalised cause-specific hazard rates are introduced, with differing degrees of heterogeneity, as summarised in Table 2.

When applying our latent class framework to describe cohort-level quantities of interest, such as the cause-specific survival and cumulative incidence functions, those quantities involving the individualised hazard rates will be expressed in terms of the variables of the latent class model; summations over the individuals in the cohort become summations over the latent classes, the relative influence of each class depending on the fraction of individuals from the cohort belonging to that class. Representing the collection of those individuals belonging to class ℓ by the variable I_ℓ, it follows that a summation over all individuals in the cohort, $\sum_{i,\mathbf{z}_i=\mathbf{z}}$, becomes a summation over the latent classes *and* their membership, of the form $\sum_\ell \sum_{i\in I_\ell,\mathbf{z}_i=\mathbf{z}}$. Those quantities expressed through summations over the individualised cause-specific hazard rates, $h_r^i(t)$, can then be written as summations over covariate-conditioned class- and cause-specific hazard rates, $h_r^\ell(t|\mathbf{z})$, as follows,

$$\frac{\sum_{i,\mathbf{z}_i=\mathbf{z}} f(h_r^i(t))}{\sum_{i,\mathbf{z}_i=\mathbf{z}} 1} = \frac{\sum_\ell \sum_{i\in I_\ell,\mathbf{z}_i=\mathbf{z}} f(h_r^\ell(t|\mathbf{z}))}{\sum_\ell \sum_{i\in I_\ell,\mathbf{z}_i=\mathbf{z}} 1} = \frac{\sum_\ell (n_\ell(\mathbf{z})/p(\mathbf{z})N) f(h_r^\ell(t|\mathbf{z}))}{\sum_\ell (n_\ell(\mathbf{z})/p(\mathbf{z})N)} = \sum_\ell w_\ell(\mathbf{z}) f(h_r^\ell(t|\mathbf{z})), \tag{7}$$

where the class membership fraction, $w_\ell(\mathbf{z})$, given by the quantity $n_\ell(\mathbf{z})/p(\mathbf{z})N$, follows from the fraction of individuals from the cohort who belong to class ℓ and have covariates $\mathbf{z}$ (where $p(\mathbf{z})$ represents the population probability of the covariate vector $\mathbf{z}_i$ being equal to $\mathbf{z}$). The class membership fractions are subject to the constraint $\sum_\ell w_\ell(\mathbf{z}) = 1$. In this paper, to avoid further complexity the class membership fractions are chosen to be independent of the covariates, this amounts to the assumption that all sub-cohorts have identically distributed covariates.

The various crude and decontaminated quantities which describe a cohort are given below in terms of our latent class parametrisation for the fully heterogeneous variant of the individualised cause-specific hazard rate ($M = 3$); corresponding expressions are easily obtained for the simpler model variants on substitution of class-independent association(s) ($M = 1$) and class-independent base hazard rate(s) ($M = 2$) as appropriate. To aid readability the following compact notation is used: the effect of the frailty and associations for each true risk r and class ℓ are summarised by the product $\boldsymbol{\beta}_r^\ell \cdot \mathbf{z} = \beta_r^{\ell 0} + \sum_\mu \beta_r^{\ell\mu} z^\mu$, the time integrals can be factorised and denoted by $\Lambda_r^\ell(t) = \int_0^t \mathrm{d}s\ \lambda_r^\ell(s)$ as the base hazard rates contain *all* time dependencies. The end-of-trial risk is assumed not to depend on the covariates, accordingly the class-independent associations for the censoring risk are defined to be zero $\boldsymbol{\beta}_0 = 0$. Summations over the covariates run over all P covariates, labelled by $\mu = 1, \ldots, P$, and shall be written as $\sum_\mu$.

In their parametrised form, the crude and decontaminated hazard rate expressions for the true risks $r = 1, \ldots, R$ are given by,

$$h_r(t|\mathbf{z}) = \frac{\sum_\ell w_\ell\ \lambda_r^\ell(t) \mathrm{e}^{\boldsymbol{\beta}_r^\ell \cdot \mathbf{z} - \sum_{r'=1}^R \exp(\boldsymbol{\beta}_{r'}^\ell \cdot \mathbf{z}) \Lambda_{r'}^\ell(t)}}{\sum_\ell w_\ell\ \mathrm{e}^{-\sum_{r'=1}^R \exp(\boldsymbol{\beta}_{r'}^\ell \cdot \mathbf{z}) \Lambda_{r'}^\ell(t)}}, \qquad \tilde{h}_r(t|\mathbf{z}) = \frac{\sum_\ell w_\ell\ \lambda_r^\ell(t) \mathrm{e}^{\boldsymbol{\beta}_r^\ell \cdot \mathbf{z} - \exp(\hat{\boldsymbol{\beta}}_r^\ell \cdot \mathbf{z}) \Lambda_r^\ell(t)}}{\sum_\ell w_\ell\ \mathrm{e}^{-\exp(\boldsymbol{\beta}_r^\ell \cdot \mathbf{z}) \Lambda_r^\ell(t)}}. \tag{8}$$

The corresponding crude and decontaminated survival functions are given by,

$$S_r(t|\mathbf{z}) = \exp\left(-\int_0^t \mathrm{d}s\ h_r(s|\mathbf{z})\right), \qquad \tilde{S}_r(t|\mathbf{z}) = \sum_\ell w_\ell\ \mathrm{e}^{-\exp(\boldsymbol{\beta}_r^\ell \cdot \mathbf{z}) \Lambda_r^\ell(t)}. \tag{9}$$

Once more, it is clear that as the crude cause-specific hazard rate is influenced by *all* risks then the crude cause-specific survival function shall be also. In Appendix A.2 the crude and decontaminated survival are shown to be identical in the case where there is one risk only, i.e. in the absence of any potential informative censoring.

The cause-specific cumulative incidence functions using our latent class model are given by,

$$F_r(t|\mathbf{z}) = \int_0^t \mathrm{d}t'\ \mathrm{e}^{-\Lambda_0(t')} \sum_\ell w_\ell\ \lambda_r^\ell(t') \mathrm{e}^{\boldsymbol{\beta}_r^\ell \cdot \mathbf{z} - \sum_{r'=1}^R \exp(\boldsymbol{\beta}_{r'}^\ell \cdot \mathbf{z}) \Lambda_{r'}^\ell(t')}. \tag{10}$$

It is noteworthy that our approach leads to an intuitive and easily interpreted formulation of the cause-specific cumulative incidence in which the role of all model parameters is completely transparent.

Our approach also offers *retrospective* determination of class membership probability for an individual using their covariates, $\mathbf{z}$, and survival information, (t, r). Following Bayesian arguments (as detailed in Appendix A.3), the probability that individual belongs to class ℓ, given their covariates and survival information, is given by,

$$P(\ell|t, r, \mathbf{z}) = \frac{w_\ell\ \lambda_r^\ell(t) \mathrm{e}^{\boldsymbol{\beta}_r^\ell \cdot \mathbf{z} - \sum_{r'=1}^R \exp\left(\boldsymbol{\beta}_{r'}^\ell \cdot \mathbf{z}\right) \Lambda_{r'}^\ell(t)}}{\sum_{\ell'=1}^L w_{\ell'}\ \lambda_r^{\ell'}(t) \mathrm{e}^{\boldsymbol{\beta}_r^{\ell'} \cdot \mathbf{z} - \sum_{r'=1}^R \exp\left(\boldsymbol{\beta}_{r'}^{\ell'} \cdot \mathbf{z}\right) \Lambda_{r'}^{\ell'}(t)}}. \tag{11}$$

The search for informative new covariates could be aided by retrospective class assignment, increasing our ability to predict personalised risk in heterogeneous cohorts. Such new covariates are expected to be features that patients in the same class have in common.

Determination of the optimal characterisation of a cohort, that is the optimal values of the latent class fractions and the frailties, associations, and base hazard rates for each class and each risk, given the available data, is described in the next section.

2.5 Characterisation of cohort heterogeneity

The Bayesian formalism has been used to determine the optimal characterisation of a cohort. The most appropriate latent class model variant, the most suitable number of latent classes, and the form of the base hazard rates are

all unknown at the outset of an analysis, so their optimal values and the corresponding optimal class membership weightings, frailties, associations, and base hazard rates all need to be determined.

The term 'model' shall be used to describe the combination of a particular number of latent classes, L, a particular form of the base hazard rates realised by a spline construction having K anchored time points, and a particular parametrisation of the personalised hazard rates, M, and shall be denoted by $\mathcal{H}_{KLM}$. The parameter vector, $\boldsymbol{\theta}_{\mathcal{H}_{KLM}}$, shall denote the model parameters (i.e. the class membership fractions, the frailty, association, and base hazard rate parameters) of the model $\mathcal{H}_{KLM}$. The expressions developed below apply for the fully heterogeneous personalised cause-specific hazard rate ($M = 3$); equivalent expressions for the simpler model variants ($M = 1$, $M = 2$) are easily obtained on substituting class-independent base hazard rate(s) and association(s) as appropriate.

The probability density to find the earliest event occurring at time t corresponding to risk r provides the link between our latent class model parameters and the observed survival data. The probability density for an individual with covariate vector $\mathbf{z}$ to report (t, r) is given by,

$$P(t, r|\mathbf{z}) = \mathrm{e}^{-\Lambda_0(t)} \sum_{\ell} w_\ell \; \lambda_r^\ell(t) \mathrm{e}^{\boldsymbol{\beta}_r^\ell \cdot \mathbf{z} - \sum_{r'=1}^{R} \exp(\boldsymbol{\beta}_{r'}^\ell \cdot \mathbf{z}) \Lambda_{r'}^\ell(t)}, \tag{12}$$

where the parameters are those familiar from the fully-heterogeneous personalised hazard rate ($M = 3$) defined in Table 2. The data likelihood, $P(D|\boldsymbol{\theta}_{\mathcal{H}_{KLM}}) = \prod_i P(t_i, r_i|\mathbf{z}_i, \boldsymbol{\theta}_{\mathcal{H}_{KLM}})$, gives the joint probability of the event times t_i and risk r_i for every individual in the cohort, conditioned on their covariates, $\mathbf{z}_i$, and the model parameters, $\boldsymbol{\theta}_{\mathcal{H}_{KLM}}$, for the model $\mathcal{H}_{KLM}$. The task of identifying the optimal characterisation of a cohort is clear; it is necessary to find the most probable parameter values, $\boldsymbol{\theta}^\star_{\mathcal{H}_{KLM}}$, of the most probable model, $\mathcal{H}^\star_{KLM}$.

The Bayesian formalism is first applied to determine the optimal parameter values, $\boldsymbol{\theta}^\star_{\mathcal{H}_{KLM}}$, for each model $\mathcal{H}_{KLM}$. The posterior distribution gives the probability of the model parameters, $\boldsymbol{\theta}_{\mathcal{H}_{KLM}}$, conditioned on the cohort survival data, D, and is given by $P(\boldsymbol{\theta}_{\mathcal{H}_{KLM}}|D) = Z^{-1}_{\mathcal{H}_{KLM}} \exp(\mathcal{L}(\boldsymbol{\theta}_{\mathcal{H}_{KLM}}, D))$, where the log-likelihood is defined as $\mathcal{L}(\boldsymbol{\theta}_{\mathcal{H}_{KLM}}, D) = \log(P(D|\boldsymbol{\theta}_{\mathcal{H}_{KLM}})P(\boldsymbol{\theta}_{\mathcal{H}_{KLM}}))$, the normalisation constant is denoted by $Z_{\mathcal{H}_{KLM}}$, and the prior distribution over the model parameters is given by $P(\boldsymbol{\theta}_{\mathcal{H}_{KLM}})$ (Appendix A.4). The contribution to the log-likelihood from the data-likelihood, for the fully heterogeneous model variant ($M = 3$), is given by,

$$\log P(D|\boldsymbol{\theta}_{\mathcal{H}_{KLM}}) = -\sum_i \Lambda_0(t_i) + \sum_i \log \left[\sum_{\ell} w_\ell \; \lambda_{r_i}^\ell(t_i) \mathrm{e}^{\boldsymbol{\beta}_{r_i}^\ell \cdot \mathbf{z}_i - \sum_{r=1}^{R} \Lambda_r^\ell(t_i) \exp(\boldsymbol{\beta}_r^\ell \cdot \mathbf{z}_i)} \right]. \tag{13}$$

In this work, maximum a posteriori parameter estimation (MAP) has been used; the optimal parameter values being those which maximise the posterior probability.

Bayesian model selection is used to determine the most probable model, $\mathcal{H}^\star_{KLM}$, given the cohort survival data, D, from an *ensemble* of models[4], where all models $\mathcal{H}_{KLM}$ are given identical prior probability $P(\mathcal{H}_{KLM})$. In pursuing this approach, the optimal model is that which is supported by the greatest "evidence" $P(\mathcal{H}_{KLM}|D)$, which is here proportional to the posterior normalisation constant $Z_{\mathcal{H}_{KLM}} = \int \mathrm{d}\boldsymbol{\theta}_{\mathcal{H}_{KLM}} \; P(D|\boldsymbol{\theta}_{\mathcal{H}_{KLM}}, \mathcal{H}_{KLM}) P(\boldsymbol{\theta}_{\mathcal{H}_{KLM}}|\mathcal{H}_{KLM})$ for that model. In this work a Gaussian approximation to the posterior has been used such that $Z_{\mathcal{H}_{KLM}}$ may be determined without the requirement for computationally expensive numerical integration (Appendix A.5), but can be determined from the optimal model parameter values, $\boldsymbol{\theta}^\star_{\mathcal{H}_{KLM}}$, and shall be written as $Z_{\mathcal{H}_{KLM}} = Z(\boldsymbol{\theta}^\star_{\mathcal{H}_{KLM}})$.

In practice, to find the optimal characterisation of a cohort, $\boldsymbol{\theta}^\star_{\mathcal{H}^\star_{KLM}}$, requires that the optimal parameter values for each of the models being compared are *first* determined and *then* the model supported by the greatest evidence is identified. The optimal parameter values, $\boldsymbol{\theta}^\star_{\mathcal{H}_{KLM}}$, are determined for every model before the optimal model, $\mathcal{H}^\star_{KLM}$, from the ensemble is determined, according to the following algorithm,

$$\mathcal{H}^\star_{KLM} = \arg\max_{\mathcal{H}_{KLM}} \left[Z(\boldsymbol{\theta}^\star_{\mathcal{H}_{KLM}}) \right], \qquad \boldsymbol{\theta}^\star_{\mathcal{H}_{KLM}} = \arg\max_{\boldsymbol{\theta}_{\mathcal{H}_{KLM}}} \left[P(D|\boldsymbol{\theta}_{\mathcal{H}_{KLM}}) P(\boldsymbol{\theta}_{\mathcal{H}_{KLM}}) \right]. \tag{14}$$

2.6 Practical tools for survival analysis

Implementation of the analysis protocol described above (Section 2.5) has been realised in our software package, *ALPACA* (**A**dvanced **L**atent Class **P**rediction **A**nd **C**ompeting Risk **A**nalysis), using the C programming language. The optimal frailty, association, and base hazard rate parameters for each model are located by MAP estimation using a stochastic refinement of the downhill simplex method [39]. Numerical estimation of the curvature of the posterior distribution around the location of maximum probability enables both the error bars for the parameter

[4] The ensemble of models is assumed to contain the actual model underlying the *true* cohort sub-structure.

estimates and the model "evidence" to be determined. As the search for the optimal parameter values is achieved using a stochastic optimization algorithm, and as the task of searching the parameter space becomes more difficult as the number of latent classes is increased and the greater the complexity of the personalised hazard rate model and base hazard rate, this procedure is typically performed multiple times and the overall optimal characterisation is selected.

A cohort's survival data is pre-processed before application of our analysis algorithms. This pre-processing involves the linear rescaling of the *raw* covariate values such that the distribution of their *transformed* counterparts correspond to a zero average and unit variance distribution, equivalent to the definition of Z-scores. Missing data values are imputed according to the average of those values which are available. The parameter estimates obtained using our analysis can be translated to the more familiar language of hazard ratios (HR), 95% confidence intervals (CI), and p-values, as a consequence of such data pre-processing. An additional benefit of this pre-processing is that the magnitudes of the effects for different associations can be directly compared.

The hazard ratio, HR_μ, associated with covariate μ, follows from generalising the concept to a balanced cohort[5], and can, as a consequence of our normalisation of covariates, be computed from the estimated association parameter β_μ according to the relation $\mathrm{HR}_\mu = \mathrm{e}^{2\beta_\mu}$. The lower and upper bounds, HR_μ^- and HR_μ^+ respectively, of the 95% confidence interval, $\mathrm{CI}_\mu = [\mathrm{HR}_\mu^-, \mathrm{HR}_\mu^+]$, for the hazard ratio associated with covariate μ are related to the estimated association parameter β_μ and its uncertainty σ_μ by the expressions $\mathrm{HR}_\mu^- = \exp(2(\beta_\mu - 1.96\sigma_\mu))$ and $\mathrm{HR}_\mu^+ = \exp(2(\beta_\mu + 1.96\sigma_\mu))$ respectively. The p-value is approximated from our parameter estimates according to the relation $p_\mu = 1 - \mathrm{erf}(|\beta_\mu|/\sqrt{2}\sigma_\mu)$ where $|\beta_\mu|$ denotes the magnitude of the estimated association parameter and the error integral is given by $\mathrm{erf}(t) = (2/\sqrt{\pi}) \int_0^t \mathrm{d}x\ \mathrm{e}^{-x^2}$.

Survival curves, cumulative incidence curves, and retrospective class allocation for each individual are all possible once the optimal frailty, association, and base hazard rate parameters have been determined. The presence of informative censoring in a cohort can be deduced from observing differences between the crude and decontaminated survival curves; the magnitude of any such differences being indicative of the extent to which competing risks mask the true cause-specific survival. False protectivity effects should be suspected if the crude survival function exceeds the decontaminated survival function; the opposite being suggestive of the influence of false exposure in the cohort data.

In a heterogeneous cohort, the class-specific decontaminated survival may offer valuable insight into the expected progression of individuals belonging to each of the latent classes. Expressing the decontaminated cause-specific survival (9), $\tilde{S}_r(t|\mathbf{z})$, as the weighted sum of cause- and class-specific survival, $\tilde{S}_r^\ell(t|\mathbf{z})$, as given by,

$$\tilde{S}_r^\ell(t|\mathbf{z}) = \mathrm{e}^{-\exp(\boldsymbol{\beta}_r^\ell \cdot \mathbf{z})\Lambda_r^\ell(t)}, \tag{15}$$

allows the class-specific decontaminated survival for each of the latent classes to be compared. Similarly, the cause-specific cumulative incidence function (10) can be expressed as a weighted sum of its class-specific components, $F_r^\ell(t|\mathbf{z})$, given by,

$$F_r^\ell(t|\mathbf{z}) = \int_0^t \mathrm{d}t'\ \mathrm{e}^{-\hat{\Lambda}_0(t')}\ \hat{\lambda}_r^\ell(t') \mathrm{e}^{\hat{\boldsymbol{\beta}}_r^\ell \cdot \mathbf{z} - \sum_{r'=1}^R \exp(\hat{\boldsymbol{\beta}}_{r'}^\ell \cdot \mathbf{z})\hat{\Lambda}_{r'}^\ell(t')}, \tag{16}$$

such that $F_r(t|\mathbf{z}) = \sum_\ell w_\ell F_r^\ell(t|\mathbf{z})$. The class- and cause-specific cumulative incidence provides information as to the relative occurrence of events for each latent class and each cause at any time during the trial.

Retrospective class assignment, by identification of the most probable latent class to which an individual belongs using (11), can offer additional insight into differences between latent classes. Retrospectively allocated class-conditioned time-to-event distributions may offer clues as to the expected survival time for members of the different classes. The detection of differences between covariate distributions for those individuals retrospectively assigned to each latent class offers a potentially powerful means of identifying novel informative covariates.

3 Application to synthetic survival data

In this section the results of the application of our algorithms to synthetic data (see Appendix B) simulating a variety of conditions are presented. The effectiveness of our algorithm to successfully characterise a heterogeneous cohort having three latent classes, two of which differ only in their base hazard rates is demonstrated in Section 3.1. In

[5] It is not possible to capture risk in a single number for non-binary covariates.

Table 3: Modelling a heterogeneous cohort: The parameters used to generate synthetic data modelling a cohort having three latent classes, one real risk, and with end-of-trial censoring at time $t = 20$. The frailty parameter, $\beta_r^{\ell 0}$, for each latent class has been set to zero.

Cohort A: ($L = 3, M = 3$) Heterogeneity and class-dependent base hazard rates			
	Class, $\ell = 1$	Class, $\ell = 2$	Class, $\ell = 3$
w_ℓ	1/3	1/3	1/3
$\lambda_1^\ell(t)$	3/10	$e^{t/4}/100$	1/10
$\beta_1^{\ell 1}$	2	2	-2
$\beta_1^{\ell 2}$	0	0	0
$\beta_1^{\ell 3}$	0	0	0

Section 3.2 the ability of our analysis to accurately characterise a cohort in the presence of heterogeneity-induced informative censoring is shown, the differences between the crude and decontaminated survival curves indicating the influence of competing risks.

The model space for the analysis of each synthetic data set covered *all* combinations of latent classes between one and four, base hazard rate complexity between one and eight, and for each of the three variants of the personalised hazard rate. The optimal parameter estimates were obtained for each model at least five times and the model with the greatest overall evidence identified to provide the optimal characterisation for each of the synthetic data sets.

As the class membership of all individuals is known when using simulated cohort data it is also possible to quantify the accuracy of Bayesian retrospective class allocation; the effectiveness of class identification using our algorithm (11) is quantified by calculating the fraction of correctly assigned individuals and the quality of allocation can be expressed for each class, q_ℓ, or for the cohort as a whole, q. It is important to note that, even if all parameters were known exactly, due to the stochasticity of event times class allocation will never be perfectly accurate[6].

3.1 Revealing heterogeneity and cohort sub-structure

Our analysis is now applied to data modelling a heterogeneous cohort having three latent classes ($L = 3$) but free of the effects of informative censoring. The characteristics of *Cohort A* are given in Table 3; the three classes are of equal size, each individual i has three covariates (z_i^1, z_i^2, z_i^3) and is subject to either one real risk or end-of-trial censoring at time $t = 20$. The personalised cause-specific hazard rates of individuals in *Cohort A* have class-dependent associations and base hazard rates (i.e. are of type $M = 3$). The associations of individuals belonging to classes $\ell = 1$ and $\ell = 2$ are identical; the classes differ only in their base hazard rates, the base hazard rate of class $\ell = 1$ is time-independent, the base hazard rate of class $\ell = 2$ increases exponentially with time. The base hazard rate of the third class $\ell = 3$ is time-independent, though it differs from that of the first class.

The characterisation of such a cohort is challenging even in the absence of informative censoring; to successfully characterise the cohort *any* analysis must be able to identify the three latent classes and distinguish between two classes which vary from each other only in their base hazard rate. The effectiveness of our analysis to characterise *Cohort A* is summarised in Figure 2, which shows the optimal estimated associations and base hazard rates for data set sizes $N = 20000$, $N = 2000$, and $N = 200$. It is evident (Figure 2) that the association parameters of *Cohort A* were accurately estimated for data set sizes $N = 2000$ and $N = 20000$, and that the estimated and true base hazard rates are sufficiently close that the optimal model can be considered as providing a good characterisation of the cohort structure. Although it is clear that a sample size of $N = 200$ is insufficient for *Cohort A* to be accurately characterised, it is noteworthy that despite such a meagre sample size our analysis correctly reports that the cohort is comprised of three latent classes and the estimated associations are accurate (within the estimated uncertainty) for covariates 2 and 3 for all three classes, and are accurate for classes 1 and 3 for covariate 1.

The optimal models reported for data set sizes $N = 20000$ and $N = 2000$ were $\mathcal{H}^\star_{KLM} = (L = 3, M = 3, K = 4)$ and $\mathcal{H}^\star_{KLM} = (K = 3, L = 3, M = 3)$ respectively. In total, 23 parameters (two weights, and for each of the three latent classes one frailty, three associations, and three modifiable spline points for the parametrised base hazard rate approximation) were estimated for the optimal model for the $N = 20000$ data set. As the Bayesian model selection stage of our analysis seeks to balance the justification for greater model complexity against the available data, it is

[6] An upper bound on the possible accuracy of retrospective allocation can be determined using the retrospective allocation and event time probability relations under the assumption of an infinitely large cohort from which the parameters have been extracted perfectly; for a cohort subject to one risk and consisting of two latent classes of equal size, with a constant class-independent base hazard rate and only one covariate, the best possible allocation quality in the case that $\beta_1^{11} = -\beta_1^{21} = 2$ is about 83%.

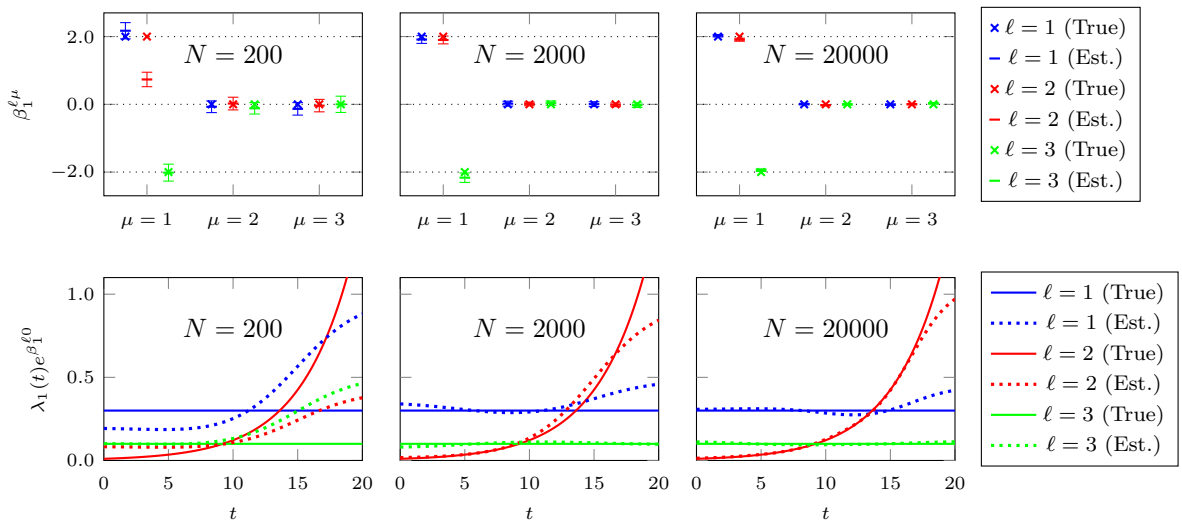

Figure 2: *Bayesian identification and characterisation of a synthetic heterogeneous cohort:* The optimal association and base hazard rate estimates for analysis of *Cohort A* data, as specified in Table 3, are shown along with their true values for sample sizes $N = 200$, $N = 2000$, and $N = 20000$. Our analysis correctly identifies three latent classes even for the meagre sample size of $N = 200$; though it is clear on inspection of the estimates that their accuracy and precision improves with increasing sample size. The association parameters are accurately estimated for all covariates and all classes on analysis of *Cohort A* data of size $N = 2000$; the base hazard rates deviate only minimally from their true value at times greater than about $t = 10$. The precision of the estimated associations is superior for the data set of size $N = 20000$ and the estimated base hazard rates are true to their actual values even beyond time $t = 15$.

unsurprising that models containing fewer parameters were found to be optimal for the $N = 2000$ and $N = 200$ data sets, having 20 and 16 parameters respectively. Ultimately, there is not enough information available in the $N = 200$ data set to justify the selection of a model having the required complexity to accurately describe *Cohort A* and the optimal model, $\mathcal{H}^{\star}_{KLM} = (K = 3, L = 3, M = 2)$, is reported as having class-independent base hazard rates.

Our analysis correctly indicates that in *Cohort A* only covariate 1 has a statistically significant association with hazard for the primary risk. The estimates for each of the three classes for the $N = 2000$ data set, for example, are (within the estimated uncertainty) in agreement with the true values ($\ell = 1$: $\beta_1^{11} = 1.90 \pm 0.12$, HR=45.96, 95% CI=[29.37,71.92], $p < 10^{-7}$; $\ell = 2$: $\beta_1^{21} = 1.91 \pm 0.11$, HR=44.75, 95% CI=[28.48,70.32], $p < 10^{-7}$; $\ell = 3$: $\beta_1^{31} = -2.16 \pm 0.14$, HR=0.013, 95% CI=[0.008,0.023], $p < 10^{-7}$). The associations for covariates 2 and 3 were correctly estimated for all of the data sets, being found to be statistically insignificant (according to p-value) in all of the three latent classes and the 95% CI included the correct HR of unity in all cases. As heterogeneity is not accounted for in a standard Cox regression it should be expected that the application of such an analysis to *Cohort A* would be likely to yield incorrect estimates. Indeed, Cox regression does produce a misleading interpretation of the *Cohort A* data having $N = 2000$, indicating that covariates 1 and 2 are both associated with an increased hazard (cov.1: HR=1.33, 95% CI=[1.21,1.47], $p < 10^{-7}$; cov. 2: HR=1.12, 95% CI=[1.01,1.23], p=0.03).

The accuracy of the spline-approximated base hazard rates can be assessed both by visual inspection, for similarity between the estimated and true rates (Figure 2), and numerically. The time-independent base hazard rate of the third class ($\ell = 3$) is reliably estimated for both data the $N = 20000$ and $N = 2000$ data sets. The base hazard rates of classes $\ell = 1$ and $\ell = 2$ are more accurately estimated for the larger data set size; the estimated base hazard rate of class $\ell = 1$ begins to deviate from the actual rate for times greater than about $t = 10$ for the $N = 2000$ analysis, whereas deviation is not observed at times earlier than about $t = 15$ for the $N = 20000$ data set. Additionally, the suitability of a spline-approximation to describe the exponentially increasing base hazard rate of the second class ($\ell = 2$) can be gauged by comparing the known values of the parameters, λ_0 and α, obtained on fitting the generalised form of the exponentially increasing base hazard rate, $\lambda(t) = \lambda_0 e^{\alpha t}$, to the estimated base hazard rate; for the $N = 20000$ data set these parameters were determined to be $\lambda_0 = 0.012$ and $\alpha = 0.233$ and for the $N = 2000$ data set they were found to be $\lambda_0 = 0.014$ and $\alpha = 0.227$ (the true values of these parameters are $\lambda_0 = 0.01$ and $\alpha = 0.25$).

The performance of retrospective class allocation will, of course, depend on the exact characteristics underlying a cohort and the accuracy of the optimal model representing that cohort; it would be reasonable to expect that retrospective class assignment should perform most effectively when the differences between latent classes are pronounced and that the algorithm may struggle when the latent classes do not differ significantly. The random assignment of patients into the three classes of *Cohort A* would offer an accuracy of 33% whereas 73% of individuals were allocated

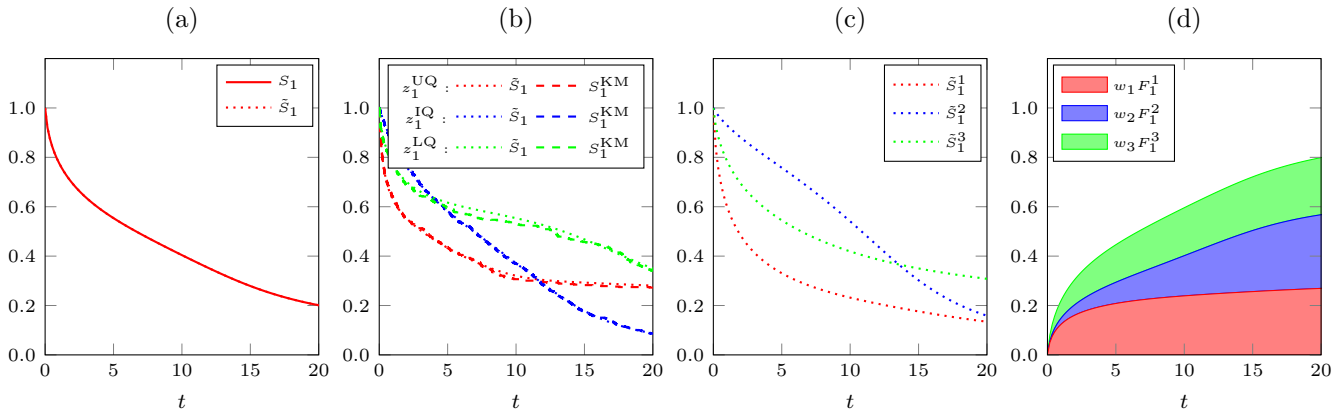

Figure 3: *Latent class-specific survival in a heterogeneous cohort:* Survival and cumulative incidence estimates arising from the optimal parameter estimates for analysis of *Cohort A* data having $N = 2000$. In (a) the crude and decontaminated survival curves for the primary risk for the cohort as a whole; as there is no informative censoring in *Cohort A* these curves are identical. In (b) the decontaminated survival curves, $\tilde{S}_1$, and the Kaplan-Meier estimators, S_1^{KM}, for the lower quartile (LQ), upper quartile (UQ), and the inter-quartile range (IQ) of covariate 1; it is clearly apparent that proportional hazards assumption does not hold for *Cohort A* as a whole. In (c) the class-specific decontaminated survival curves, $\tilde{S}_1^{\ell}$, for the three classes; the impact of the constant base hazard rate of class $\ell = 1$ and the exponentially increasing base hazard rate of class $\ell = 2$ is easily observed in the difference between the survival curves for the two classes. In (d) the stacked weighted class-specific cumulative incidence.

correctly on applying our class assignment algorithm to the $N = 20000$ data set; 72% and 66% of those individuals allocated to the first and second classes ($\ell = 1$, $\ell = 2$) actually belonged to those classes, and 84% of individuals allocated to the third class ($\ell = 3$) actually belonged to that class. Class assignment applied to the $N = 2000$ data set yielded an overall accuracy of 72% and with a similar accuracy for those assigned to the first and second class as was obtained with the $N = 20000$ data set; 78% of those allocated to the third class actually belonged to that class. It is interesting to note that for the $N = 200$ data set, despite there being insufficient information for all associations and the base hazard rates to be accurately estimated, overall 70% of individuals were correctly allocated to their true class.

In Figure 3 class-specific survival and weighted cumulative incidence functions generated using the optimal parameter estimates for the $N = 2000$ data set are shown. As there is only one real risk in *Cohort A* the crude and decontaminated survival functions, S_1 and $\tilde{S}_1$, are identical (Figure 3 (a)). As there is no informative censoring in *Cohort A* the Kaplan-Meier estimators offer a reasonable approximation to the risk-specific survival function (Figure 3 (b)); the crossing of the Kaplan-Meier estimators conditioned on the value of covariate 1 would suggest that there is latent heterogeneity in *Cohort A* and that the application of standard Cox analysis would yield misleading estimates. The differences in the expected survival between the three classes over the duration of the trial is readily apparent on inspecting the class-specific decontaminated survival functions (Figure 3 (c)); the significantly more rapid expiration of the first class ($\ell = 1$) against the primary risk survival in comparison to that of the second class ($\ell = 2$) being due to the two classes having very different base hazard rates. The weighted class-specific cumulative incidence (Figure 3 (d)) also illustrate at which stages of the trial members of the first ($\ell = 1$) and second class ($\ell = 2$) succumb to the primary risk.

3.2 Effective analysis in the presence of informative censoring

The ability of our analysis to accurately model the true survival in the presence of heterogeneity-induced informative censoring is now examined. The analysis of two cohorts having identical primary risk characteristics are presented in this section; one cohort being modelled to include the effects of false protectivity and the other to include the influence of false-aetiology.

Both *Cohort B* and *Cohort C* have two latent classes ($L = 2$) of equal size and have personalised cause-specific hazard rates having class-dependent associations (i.e. of type $M = 2$); both are subject to two real risks, the primary risk ($r = 1$) and a competing risk ($r = 2$), and end-of-trial censoring. The base hazard rates are time-independent but differ between the two risks. The properties of *Cohort B* and *Cohort C* are given in Table 4 alongside the Bayesian-determined optimal parameter values for data having $N = 1500$ observations.

The optimal model was correctly determined for the analysis of both *Cohort A* and *Cohort B* data, describing heterogeneous cohorts containing two classes and with personalised hazard rates having heterogeneous associations but with class-independent and time-independent base hazard rates. The predicted class sizes, associations, and base

Table 4: Modelling heterogeneity-induced informative censoring: The estimated weights, w_ℓ, and association parameters, $\beta_r^{\ell\mu}$, for the two real risks $r = 1$ and $r = 2$, from the Bayesian-determined optimal model from analysis of the *Cohort B* and *Cohort C* data sets with $N = 1500$. The true parameter values are shown in brackets to the right of the estimated values.

Cohort B $(K = 1, L = 2, M = 2)$ Heterogeneity-induced false protectivity				
	Class, $\ell = 1$		Class, $\ell = 2$	
w_ℓ	0.51 ± 0.02	(0.5)	0.49 ± 0.02	(0.5)
$\beta_1^{\ell 1}$	1.85 ± 0.14	(2.0)	-1.97 ± 0.10	(-2.0)
$\beta_1^{\ell 2}$	-0.05 ± 0.10	(0.0)	-0.04 ± 0.07	(0.0)
$\beta_1^{\ell 3}$	0.15 ± 0.10	(0.0)	0.01 ± 0.09	(0.0)
$\beta_2^{\ell 1}$	3.17 ± 0.10	(3.0)	-0.10 ± 0.07	(0.0)
$\beta_2^{\ell 2}$	-0.07 ± 0.08	(0.0)	0.03 ± 0.06	(0.0)
$\beta_2^{\ell 3}$	-0.06 ± 0.07	(0.0)	0.01 ± 0.06	(0.0)

Cohort C $(K = 1, L = 2, M = 2)$ Heterogeneity-induced false aetiology				
	Class, $\ell = 1$		Class, $\ell = 2$	
w_ℓ	0.51 ± 0.02	(0.5)	0.49 ± 0.02	(0.5)
$\beta_1^{\ell 1}$	2.04 ± 0.12	(2.0)	-1.94 ± 0.10	(-2.0)
$\beta_1^{\ell 2}$	-0.06 ± 0.07	(0.0)	-0.01 ± 0.06	(0.0)
$\beta_1^{\ell 3}$	-0.02 ± 0.08	(0.0)	0.14 ± 0.06	(0.0)
$\beta_2^{\ell 1}$	-3.02 ± 0.09	(-3.0)	0.11 ± 0.08	(0.0)
$\beta_2^{\ell 2}$	-0.02 ± 0.06	(0.0)	0.05 ± 0.07	(0.0)
$\beta_2^{\ell 3}$	0.04 ± 0.07	(0.0)	-0.03 ± 0.06	(0.0)

hazard rates for both data sets were in close agreement with their actual values, as detailed in Table 4.

The Bayesian-determined estimates for both *Cohort B* and *Cohort C* suggest correctly that covariate 1 is strongly associated with an increased hazard for the primary risk in one class and is associated with a reduced hazard in the other class; covariates 2 and 3 were found to be statistically insignificant, according to their respective p-values, for both classes in both cohorts.

The Bayesian-determined estimates for the secondary risk, responsible for informative censoring, were also found to be consistent with the expected values for both cohorts. In *Cohort B* covariate 1 was found to be strongly associated with an increased hazard for the secondary risk in one class only ($\ell = 1$: HR=565.28, 95% CI=[387.72,824.17], $p = 10^{-7}$), and in *Cohort C* covariate 1 was found to be associated with a reduced hazard for the secondary risk for one class only ($\ell = 1$: HR=0.002, 95% CI=[0.002,0.003], $p = 10^{-7}$). Again, the associations for the other covariates were found to be statistically insignificant and the predicted hazard ratios suggested neither elevated or reduced hazard for the secondary risk.

It is reassuring that our analysis was able to correctly characterise the associations for the primary risk as being identical (within the estimated uncertainty of the parameter estimates) both in the presence of heterogeneity-induced false protectivity (*Cohort B*) and in the presence of heterogeneity-induced false aetiology (*Cohort C*). As was found in Section 3.1 for the analysis of heterogeneous cohort data, standard Cox regression offered misleading estimates on the analysis of *Cohort B* and *Cohort C* data.

The various survival estimators obtained with the optimal parameter estimates for *Cohort B* are shown in Figure 4. False protectivity in the cohort is clearly indicated (Figure 4 (a)); the crude survival function for the primary risk, S_1, exceeds its decontaminated counterpart, $\tilde{S}_1$, consequently *true* survival against the primary risk is poorer than the crude survival function suggests. In the presence of informative censoring the Kaplan-Meier estimators for the primary risk, S_1^{KM}, conditioned on the value of covariate 1, are also seen to be misleading (Figure 4 (b)); the risk-specific Kaplan-Meier estimators follow the crude survival function (not shown) and yield extremely poor estimates for the lower and upper quartiles of covariate 1 due to risk correlations in the cohort. The Kaplan-Meier estimator, for example, indicates that individuals having an upper quartile value of covariate 1 are likely to have a relatively good survival (about 90% surviving at the end-of-trial time) against the primary risk whereas such individuals actually have a much poorer survival (less than 50% survival at the end-of-trial time). The class-specific survival functions and weighted class-specific cumulative incidence are also shown (Figure 4 (c,d)).

In Figure 5 the survival estimators for *Cohort C* are shown. The influence of false aetiology effects in *Cohort C* is clear on noting that the decontaminated survival function for the primary risk, $\tilde{S}_1$, exceeds the crude survival function, S_1, throughout the trial duration; the *true* survial against the primary risk is significantly greater than the crude survival function suggests. The covariate-conditioned risk-specific Kaplan-Meier estimators prove also to be very misleading in the presence of false exposure effects in *Cohort C*; the Kaplan-Meier estimator for individuals having a lower quartile value of covariate 1 suggests no survival beyond about half of the trial duration whereas the survival for such individuals is actually almost 50% at the end-of-trial time. The class-specific survival functions and weighted class-specific cumulative incidence are also shown (Figure 5 (c,d)).

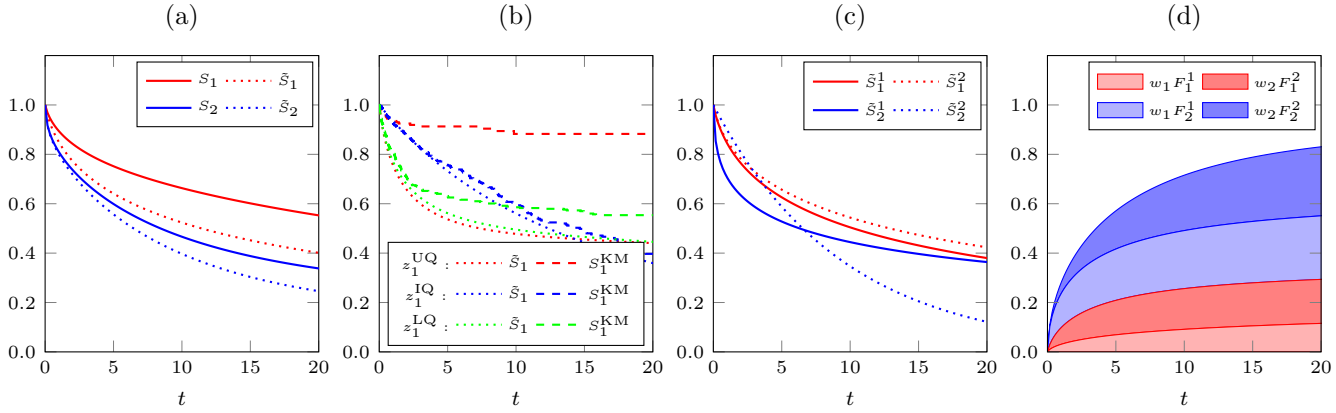

Figure 4: *Decontaminating false protectivity effects from survival data:* Survival and cumulative incidence curves for analysis of *Cohort B* (Table 4) data having $N = 1500$ demonstrating the effectiveness of our analysis to expose the influence of false protectivity in survival data. In (a) the crude and decontaminated risk-specific survival curves, $S_r(t)$ and $\tilde{S}_r(t)$, for the primary and secondary risks; the extent to which the secondary risk causes false protectivity is clearly apparent on observing the difference between the crude and decontaminated survival curves for the primary risk. In (b) the decontaminated survival curves, $\tilde{S}_1$, and the risk-specific Kaplan-Meier estimators, S_1^{KM}, for the lower quartile (LQ), upper quartile (UQ), and the inter-quartile range (IQ) of covariate 1; it is clear that proportional hazards assumption does not hold for *Cohort B* as a whole and it is striking that the risk-specific Kaplan-Meier estimator for the upper quartile of covariate 1 is extremely misleading. In (c) the class- and risk-specific decontaminated survival curves, $\tilde{S}_r^{\ell}$, for the two latent classes for both risks. In (d) the stacked weighted class-specific cumulative incidence.

4 Applications to prostate cancer data

Prostate cancer (PC) data are notorious for exhibiting competing risk effects [40], largely due to the fact that the disease occurs late in life when there is an increased number of non-primary events whose incidence could correlate with prostate cancer. Here we analyse data from the ULSAM cohort [35] and compare the outcomes of Cox's proportional hazards regression [11] and our present method. The ULSAM cohort has $N = 2047$ individuals of which 208 reported PC as the first event as described previously [36] and we have included five relevant covariates of the ULSAM data [35].[7]

Smoking is suggested to have a weak protective effect against PC risk (HR=0.85, 95% CI=[0.65,1.11], p=0.23) in the ULSAM cohort according to Cox's proportional hazards analysis. However, as the Kaplan-Meier estimator for PC risk conditioned on the smoking covariate (see Figure 1) does not meet the proportional hazards assumption, any conclusions drawn from such an analysis (on the ULSAM cohort data as a whole, at least) may be invalid. The non-proportionality of the smoking covariate-conditioned Kaplan-Meier estimator for PC risk could be due to heterogeneity in the ULSAM cohort.

The optimal characterisation of the ULSAM data, as determined by Bayesian model selection, suggests that the cohort should be viewed as consisting of two distinct classes: one class $\ell = 1$ with relatively frail individuals (in terms of both primary and secondary risk) which contains about 16% of the cohort according to retrospective class allocation, and another class $\ell = 2$ with rather healthy individuals which contains the remainder of the cohort. The estimated association parameters, base hazard rates, and class-specific survival functions for PC are shown in Fig. 6 for the most probable model ($\mathcal{H}^{\star}_{KLM} = (K = 3, L = 2, M = 2)$). The decontaminated survival function for PC was found to be marginally less than the crude survival function for follow-up times exceeding about 20 years; this is indicative of false protectivity effects for PC risk in the ULSAM cohort data.

In the Bayesian-determined two-class description of the ULSAM cohort, BMI and smoking are recognised as serious PC risk factors for those individuals in the healthier class. In the healthier class ($\ell = 2$) smoking is associated with elevated risk of PC (HR=4.08, 95% CI=[1.09,15.29], p=0.04), as is having a higher BMI (HR=2.77, 95% CI=[1.20,6.36], p=0.02). Conversely, in the frailer class ($\ell = 1$) smoking is associated with a decreased risk of PC (HR=0.30, 95% CI=[0.12,0.76], p=0.01). In the frail class the regression coefficients are weaker than those of the stronger class, and one expects the negative coefficients for e.g. BMI and smoking to reflect reverse causality: within this group, having a higher BMI and still being *able* to smoke may well be an indicator of *relatively* good health.

Our explanation of the ULSAM data is not necessarily the final one. There are alternative ways to do the regression, e.g. by combining all non-primary risks (including the end-of-trial risk). The conclusion to be drawn is that the new two-class explanation of the ULSAM data is both probabilistically and intuitively more plausible.

[7] Since this study does not seek to answer a pre-specified clinical hypothesis, the actual choice made for the set of covariates to be included is not critical.

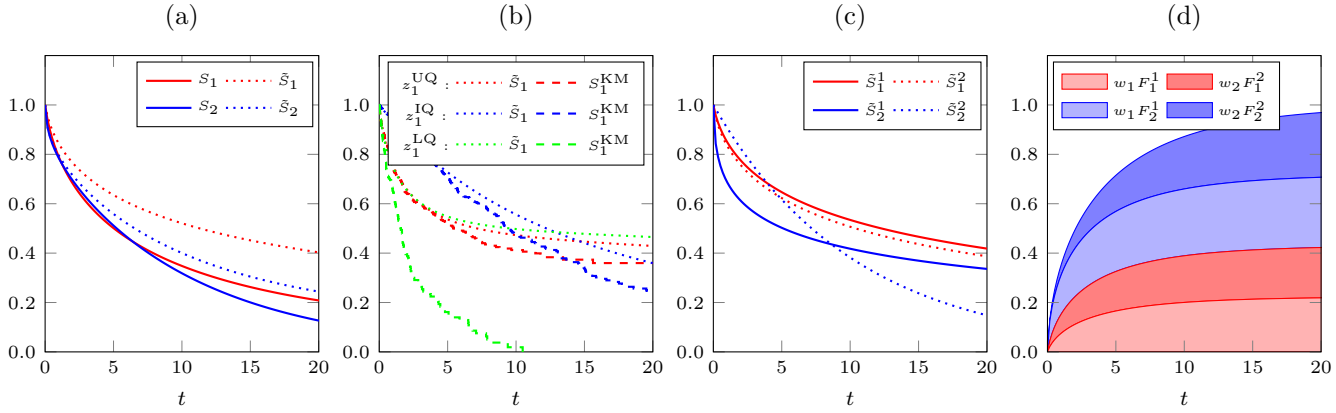

Figure 5: *Decontaminating false aetiology effects from survival data:* Survival and cumulative incidence curves for analysis of *Cohort C* (Table 4) data having $N = 1500$ demonstrating the effectiveness of our analysis to expose the influence of false aetiology in survival data. In (a) the crude and decontaminated risk-specific survival curves, $S_r(t)$ and $\tilde{S}_r(t)$, for the primary and secondary risks; the extent to which the secondary risk suggests false exposure is clearly apparent on observing the difference between the crude and decontaminated survival curves for the primary risk. In (b) the decontaminated survival curves, $\tilde{S}_1$, and the risk-specific Kaplan-Meier estimators, S_1^{KM}, for the lower quartile (LQ), upper quartile (UQ), and the inter-quartile range (IQ) of covariate 1; it is clear that proportional hazards assumption does not hold for *Cohort C* as a whole and that the risk-specific Kaplan-Meier estimator for the lower quartile of covariate 1 is extremely misleading. In (c) the class- and risk-specific decontaminated survival curves, $\tilde{S}_r^\ell$, for the two latent classes for both risks. In (d) the stacked weighted class-specific cumulative incidence.

5 Applications to breast cancer data

In this section the results of the application of our analysis to data from the Swedish Apolipoprotein Mortality Risk Study (AMORIS) e.g. [37, 38] are summarised. The analysed data set is that for which a competing risk analysis is presented in [41], and contained data for $N = 1798$ women from the AMORIS population for whom baseline serum glucose, triglyceride, and total cholesterol measurements were available within three months to three years prior to breast cancer diagnosis. The time-to-event describes the duration between diagnosis and breast cancer death (BC), cardiovascular disease death (CV), death from other causes, or departure from or the end of the study (censoring) and age, fasting status, and socio-economic status data were also included as covariates.

Our analysis suggests that the cohort of women from the AMORIS population diagnosed with breast cancer is best described by three latent classes which share the same general trend of having base hazard rates which decrease with time for BC death risk but increase with time for CV death risk. The largest class ($\ell = 1$) contains about 66% of the cohort according to retrospective class allocation and has a poorer survival against BC death but a greater survival against CV death than does the second class ($\ell = 2$), which accounts for 32% of the cohort. The smallest latent class ($\ell = 3$) is almost wholly comprised of relatively younger individuals for whom death from other causes was reported; it is likely that the inclusion of this third class, containing only 35 of the 1798 individuals according to retrospective class assignment, enables more effective characterisation of the other two classes as it allows the effects of two differing groups within those reported as having death from other causes to be distinguished. As BC and CV death shall be the focus of the remainder of this section, parameter estimates and survival curves pertaining to the small third latent class ($\ell = 3$) are omitted in the interest of readability. The estimated associations, base hazard rates, and decontaminated class-specific survival curves are shown in Figure 7 for the optimal characterisation of this cohort.

The largest class ($\ell = 1$) has the greatest hazard for BC death, having a base hazard rate which exceeds that of the second class ($\ell = 2$) over the entire interval (Figure 7 (c)); the opposite is true for the hazard for CV death, with the second class ($\ell = 2$) having the greatest hazard for CV death. The base hazard rate for BC death is most significant at diagnosis ($t = 0$) and diminishes with time whereas the risk of CV death increases with time. The base hazard for CV death exceeds the risk of BC death after about 20 years from diagnosis for those in the first class ($\ell = 1$) and after about 16 years for those in the second class ($\ell = 2$); this is generally consistent with previous findings regarding the importance of considering co-morbidities for women with breast cancer [42]. Accordingly, survival against BC death is greater for the less frail second class ($\ell = 2$) than for the first class ($\ell = 1$) and survival against CV death is greater in the first class ($\ell = 1$) than in the second class ($\ell = 2$).

In the relatively frail first class ($\ell = 1$), triglyceride levels were found to be associated with an increased hazard for both CV death (CV, $\ell = 1$: HR=6.44, 95% CI=[1.86,22.29], $p = 0.003$) and for BC death (BC, $\ell = 1$: HR=1.62, 95% CI=[0.92,2.84], $p = 0.09$). Standard Cox analysis suggests a weaker association with an increased hazard against CV

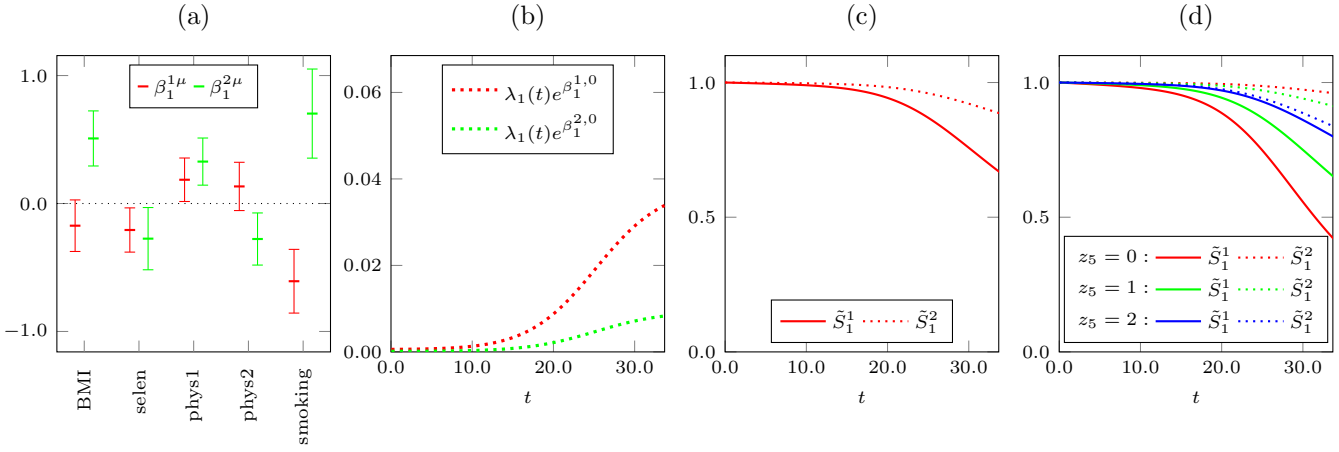

Figure 6: *Bayesian-determined heterogeneous characterisation of the ULSAM cohort:* In (a) and (b) the associations and frailty weighted-base hazard rates for PC risk, for the two latent classes of the optimal model $\mathcal{H}^{\star}_{KLM} = (K = 3, L = 2, M = 2)$ obtained on analysis of the ULSAM cohort data. The included covariates are: body mass index (real-valued), serum Selenium level (*selen*, integer valued), leisure time physical activity (*phys1*, discrete levels 0/1/2), work physical activity (*phys2*, discrete levels 0/1/2), and smoking status (*smoking*, discrete levels 0/1/2). The estimated weighting of the classes for this model is $w_1 = 0.32 \pm 0.08$, $w_2 = 0.68 \pm 0.08$. The retrospective weighting of the classes for this model is $f_1 = 0.16$ (332 of 2047 patients), $f_2 = 0.84$ (1715 of 2047 patients). In (c) the class-and risk-specific survival functions; for PC risk it is clear that the survival of members of the frailer class ($\ell = 1$) is poorer than those of the healthier class ($\ell = 2$), significantly and increasingly so for times beyond about 20 years. In (d) the class-and risk-specific survival function conditioned on the three values of the smoking covariate ($z_5 = 0$: non-smoker, $z_5 = 1$: ex-smoker, $z_5 = 2$: smoker); it is clear that for the majority of the cohort ($\ell = 2$) smoking is associated with poorer survival outcomes but that for those members of the frailer class ($\ell = 1$) smokers and ex-smokers have greater survival than non-smokers.

death (CV, Cox: HR=1.61, 95% CI=[1.17,2.21], $p = 0.003$) than is suggested by our analysis, a probable consequence of those members of the cohort for which triglyceride levels are not associated with increased hazard ($\ell = 2$) diluting the effects of those for which triglyceride levels are associated with increased hazard ($\ell = 1$); the same is also true for BC death, for which standard Cox analysis indicates that triglyceride levels are associated with a modestly increased hazard for BC death (BC, Cox: HR=1.22, 95% CI [0.98,1.52], p=0.076). Serum glucose was found to be associated with an increased hazard for CV death (CV, $\ell = 2$: HR=1.59, 95% CI=[1.01,2.49], $p = 0.04$) for those in the relatively less frail second class ($\ell = 2$), as was total cholesterol (CV, $\ell = 2$: HR=1.65, 95% CI=[0.79,3.51], $p = 0.18$).

In contrast to standard Cox analysis, which suggests that age is associated with a reduced hazard for BC death (BC, Cox: HR=0.77, 95% CI=[0.59,1.00], p=0.05) for the cohort as a whole, our analysis indicates that for the majority of the cohort ($\ell = 1$) age is associated with an increased hazard for BC death (BC, $\ell = 1$: HR=3.49, 95% CI=[1.51,8.10], p=004) and is associated with a reduced hazard for BC death (BC, $\ell = 2$: HR=0.09, 95% CI=[0.02,0.30], $p \approx 10^{-4}$) for about only one third of the cohort ($\ell = 2$). The age of an individual was also found to be associated with their hazard for CV death, most significantly being strongly associated with an increased hazard for CV death (CV, $\ell = 2$: HR=42.93, 95% CI=[16.43,112.21], $p < 10^{-7}$) in the less frail second class ($\ell = 2$). The age-conditioned risk-specific Kaplan-Meier estimator and class-specific decontaminated survival for BC death are shown in Figure 8, along with the age, time-to-event, and retrospectively assigned class, for those individuals recorded as succumbing to BC death. The stratification by age, between about 50 and 55 years, of those individuals succumbing to BC death and retrospectively allocated to the relatively frail class ($\ell = 1$) and those allocated to the second class ($\ell = 2$), suggest that menopausal status may be informative as to expected survival against BC death, as shown in Figure 8(c).

Survival against BC death for the more frail first class ($\ell = 1$) is greatest for the younger members of the cohort, with over 80% survival for those in the LQ for age in the cohort as opposed to less than 50% survival for those in the UQ for age. In the second class ($\ell = 2$) survival against BC death is markedly poorer for younger members of the cohort as a consequence of age being strongly associated with reduced hazard for BC death for this class.

6 Discussion

In this study we have introduced a fully Bayesian approach for the determination of the probabilistically optimal characterisation of a complex cohort. The introduction of the concept of heterogeneity-induced informative censoring, where risks are independent only at the level of individuals, has enabled us to identify the 'decontaminated' hazard rates and survival functions.

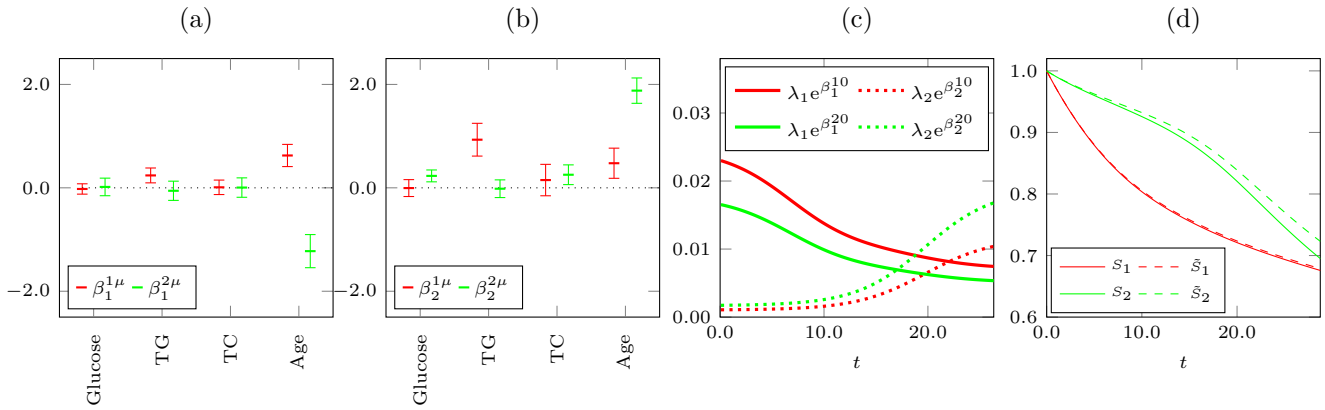

Figure 7: *Bayesian-determined heterogeneous characterisation of a cohort of women from the AMORIS population diagnosed with breast cancer:* Parameter estimates for the Bayesian-determined optimal model, $\mathcal{H}^{\star}_{KLM} = (K = 3, L = 3, M = 2)$, obtained on analysis of data for $N = 1798$ women from the AMORIS population diagnosed with breast cancer; the estimated retrospective weighting of the classes for this model is $f_1 = 0.66$ (1189 of 1798 patients), $f_2 = 0.32$ (574 of 1798 patients), $f_3 = 0.02$ (35 of 1798 patients). In (a) and (b) the estimated hazard associations, $\beta_r^{\ell\mu}$, with serum glucose, triglyceride (TG), total cholesterol (TC) and age, for BC death ($r = 1$) and CV death ($r = 2$) respectively for the dominant ($\ell = 1$: red) and next largest ($\ell = 2$: green) latent classes. In (c) the base hazard rates for BC and CV death for the two dominant latent classes; note that the base hazard rate for CV death exceeds that for BC death for follow-up times greater than about 15 years ($\ell = 1$: red, $\ell = 2$: green). In (d) the crude (solid lines) and decontaminated (dashed lines) cause-specific survival against BC (red lines) and CV death (green lines); observe that the crude survival against CV death, S_2, is poorer than its decontaminated counterpart, $\tilde{S}_2$.

The software package, *ALPACA*, which implements the formalism introduced herein provides practical tools for survival analysis which can be applied to heterogeneous cohorts and in the presence of heterogeneity-induced informative censoring. This implementation is able to find the most probable characterisation of a cohort through the combination of i) a latent class model which captures, for *all* risks and for *each* class, the base hazard rate, covariate association(s) and relative frailty, and ii) Bayesian model selection to determine the optimal number of classes, the optimal parametrisation of each base hazard rate, and the extent of heterogeneity in the cohort. Once the optimal description of the cohort has been determined, crude and decontaminated cause-specific survival curves can be compared to gauge the extent of any informative censoring, and differences between class-specific survival curves can be examined. Exploration of correlations between covariate values and retrospectively assigned class membership can offer an extra insight into a cohort and may aid the search for new informative biomarkers. An additional advantage of our approach is that through the incorporation of Bayesian model selection to determine the optimal description of a cohort, our survival analysis can be applied with little interaction or effort required from the analyst prior to inspection and assessment of the frailties, covariate association(s), and base hazard rates for the optimal model.

Applied to synthetic data, our analysis was shown to effectively characterise heterogeneous cohorts, successfully remove heterogeneity-induced false protectivity and false aetiology effects, and even discriminate between two classes differing only in their base hazard rates. Retrospective class allocation was demonstrated to have an impressive accuracy even for survival data modelling a cohort containing 200 individuals.

On application to real survival data from the ULSAM cohort, with prostate cancer as the primary risk, our analysis leads to plausible alternative explanations for previous counter-intuitive inferences (such as a weak protective effect on PC of smoking), in terms of distinct sub-groups of patients with distinct risk factors and overall frailties. In the ULSAM cohort, it was also shown that the men's metabolic status introduce competing risk problems, although using traditional methods, several different analyses had to be done to reveal the underlying risk pattern [40], which with our proposed method could be done coherently in one analysis.

Applied to survival data for women diagnosed with breast cancer from the Swedish Apolipoprotein Mortality Risk Study (AMORIS), our analysis suggests that for breast cancer death risk the base hazard rate decreases with time but increases with time for cardiovascular death risk, and the age-class membership correlations may suggest differences in association patterns and survival against breast cancer between pre- and post-menopausal women.

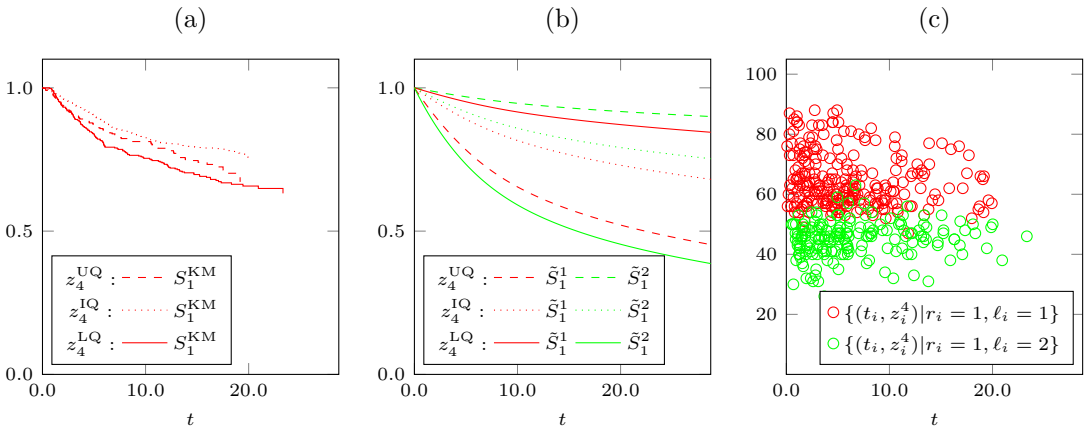

Figure 8: *Age and survival in women from the AMORIS population diagnosed with breast cancer:* In (a) the age-conditioned Kaplan-Meier estimator for BC death; observe that survival of both the upper and lower quartiles of age the proportional hazards is poorer than that for the inter-quartile range for age (i.e. the proportional hazards assumption is not met) and is suggestive of heterogeneity in the cohort. In (b) the class-specific decontaminated survival curves, $\tilde{S}_r^\ell$ conditioned on the age covariate (z_4), for BC death ($r = 1$). The substantial difference in survival against BC death between those in the lower quartile (LQ) and upper quartile (UQ) for age is indicative of the strength of the association of age with hazard for BC death (see Figure 7). Observe that in the more frail first class ($\ell = 1$) survival against BC death decreases with increasing age while the opposite is true of the second class ($\ell = 2$); however, survival of the second class ($\ell = 2$) is actually poorer for younger individuals than it is for older from the more frail first class ($\ell = 1$). In (c) the time-to-event and age for each individual having succumbed to BC death is shown, indicating that there is a clear stratification, according to age, of the latent classes to which individuals are retrospectively assigned. Almost all of those individuals younger than about 50 years old are assigned to the second latent class ($\ell = 2$) whereas those older than about 50 years are assigned to the first class ($\ell = 1$).

References

[1] Tsiatis A. A nonidentifiability aspect of the problem of competing risks. *Proceedings of the National Academy of Sciences USA* 1975; **72**: 20-22.

[2] Gail M. A review and critique of some models used in competing risk analysis. *Biometrics* 1975; **31**: 209-222.

[3] Klein JP. Competing risks. *Wiley Interdisciplinary Reviews: Computational Statistics* 2010; **2**: 333-339.

[4] Andersen PK, Geskus RB, de Witte T, Putter H. Competing risks in epidemiology: possibilities and pitfalls. *International Journal of Epidemiology* 2012; **41**: 861-870.

[5] Di Serio C. The protective impact of a covariate on competing failures with an example from a bone marrow transplantation study. *Lifetime Data Analysis* 1997; **3**: 99-122.

[6] Scharfstein DO, Robins JM. Estimation of the failure time distribution in the presence of informative censoring. *Biometrika* 2002; **89**: 617-634.

[7] Dignam JJ, Zhang Q, Kocherginsky M. The use and interpretation of competing risk regression models. *Clinical Cancer Research* 2012; **18**: 2301-2306.

[8] Thompson CA, Zhang Z-F, Arah OA. Competing risk bias to explain the inverse relationship between smoking and malignant melanoma. *European Journal of Epidemiology* 2013; **28**: 557-567.

[9] Soneji S, Beltrán-Sánchez H, Sox HC. Assessing progress in reducing the burden of cancer mortality, 1985-2005. *Journal of Clinical Oncology* 2014; **32(5)**: 444-448.

[10] Kaplan EL, Meier P. Nonparametric estimation from incomplete observations. *Journal of the American Statistical Association* 1958; **53**: 457-481.

[11] Cox DR. Regression models and life tables. *Journal of the Royal Statistical Society* B 1972; **34**: 187-220.

[12] Kalbfleisch DL, Prentice RL. *The statistical analysis of failure time data*. New York, Wiley: 2002.

[13] Pintilie M. *Competing Risks: A Practical Perspective*. Wiley: 2006.

[14] Lancaster T. Econometric methods for the duration of unemployment. *Econometrica* 1979; **47**: 939-956.

[15] Vaupel JW, Manton KG, Stallard E. The impact of heterogeneity in individual frailty on the dynamics of mortality. *Demography* 1979; **16**: 439-454.

[16] Zahl PH. Frailty modelling for the excess hazard. *Statistics in Medicine* 1997; **16**: 1573-1585.

[17] Yashin AI, Iachine IA. Genetic analysis of durations: correlated frailty model applied to survival of danish twins. *Genetic Epidemiology* 2005; **12**: 529 -538.

[18] Gorfine M, Hsu L. Frailty-based competing risk model for multivariate survival data. *Biometrics* 2011; **67**: 415426.

[19] Keiding N, Andersen PK, Klein JP. The role of frailty models and accelerated failure time models in describing heterogeneity due to omitted covariates. *Statistics in Medicine* 1997; **16**: 215-224.

[20] Vaida F, Xu R. Proportional hazard model with random effects. *Statistics in Mededicine* 2000; **19**: 3309-3324.

[21] Rosner B, Glynn RJ, Tamimi RM, Chen WY, Colditz GA, Willett WC, Hankinson SE. Breast cancer risk prediction with heterogeneous risk profiles according to breast cancer tumour markers. *American Journal of Epidemiology* 2013; **178**: 296-308.

[22] Wienke A. *Frailty Models in Survival Analysis*. Chapman & Hall/CRC Biostatistics Series: Boca Raton, 2010.

[23] Duchateau L, Jansen P. *The Frailty Model (Statistics for Biology and Health)*. Springer: New York, 2008.

[24] Lazarsfeld PF. The logical and mathematical foundation of latent structure analysis. *SA Stouffer at al. (Eds) Measurement and Prediction*. Princeton: Princeton University Press, 1950.

[25] Huang X, Wolfe RA. A frailty model for informative censoring. *Biometrics* 2002; **58**: 510-520.

[26] Muhten B, Masyn K. Discrete-time survival mixture analysis. *Journal of Educational and Behavioral Statistics* 2005; **30**: 27-58.

[27] Reboussin BA, Anthony JC. Latent class marginal regression models for modelling youthful drug involvement and its suspected influences. *Statistics in Medicine* 2001; **20**: 623-639.

[28] Fine JP, Gray RJ. A proportional hazards model for the subdistribution of a competing risk. *Journal of the American Statistical Society* 1999; **94**: 496-509.

[29] Fine JP. Regression modeling of competing crude failure probabilities. *Biostatistics* 2001; **2**: 85-97.

[30] Klein JP, Andersen PK. Regression modeling of competing risks data based on pseudovalues of the cumulative incidence function. *Biometrics* 2005; **61**: 223-229.

[31] Jeong JH, Fine JP. Parametric regression on cumulative incidence function. *Biostatistics* 2007; **8**: 184-196.

[32] Katsahian S, Boudreau C. Estimating and testing for center effects in competing risks. *Statistics in Medicine* 2011; **30**: 1608-1617.

[33] Heckman JJ, Honoré BE. The identifiability of the competing risks model. *Biometrika* 1989; **76**: 325-330.

[34] Abbring JH, van den Berg GJ. The identifiability of the mixed proportional hazards competing risks model. *Journal of the Royal Statistical Society B* 2003; **65**: 701-710.

[35] *Uppsala Longitudinal Study of Adult Men.* Department of Public Health and Caring Sciences/ Geriatrics, Uppsala Universitet. www2.pubcare.uu.se/ULSAM/.

[36] Grundmark B, Zethelius B, Garmo H, Holmberg L. Serum levels of selenium and smoking habits at age 50 influence long term prostate cancer risk; a 34 year ULSAM follow-up. *BMC Cancer* 2011; **11**: 431.

[37] Holme I, Aastveit AH, Hammar N, Jungner I, Walldius G. Inflammatory markers, lipoprotein components and risk of major cardiovascular events in 65,005 men and women in the Apolipoprotein MOrtality RISk study (AMORIS). *Atherosclerosis.* 2010; **213(1)**: 299-305.

[38] Holme I, Aastveit AH, Jungner I, Walldius G. Relationships between lipoprotein components and risk of myocardial infarction: age, gender and short versus longer follow-up periods in the Apolipoprotein MOrtality RISk study (AMORIS). *J. Intern. Med.* 2008; **264(1)**: 299-305.

[39] Press WH, Teukolsky SA, Vetterling WT, Flannery BP. *Numerical recipes in C - the art of scientific computing.* Cambridge: Cambridge University Press, 1992.

[40] Grundmark B, Garmo H, Loda M, Busch C, Holmberg L, Zethelius B. The metabolic syndrome and the risk of prostate cancer under competing risks of death from other causes. *Cancer Epidemiol Biomarkers Prev.* 2010; **19**: 2088-2096.

[41] Wulaningsih W, Vahdaninia M, Rowley M, Holmberg L, Garmo H, Malmstrom H, Lambe M, Hammar N, Walldius G, Jungner I, Coolen ACC, Van Hemelrijck M. Prediagnostic serum glucose and lipids in relation to survival in breast cancer patients: a competing risk analysis. *BMC Cancer* (in press).

[42] Patnaik JL, Byers T, Di Guiseppi C, Dabelea D, Denberg TD. Cardiovascular disease competes with breast cancer as the leading cause of death for older females diagnosed with breast cancer: a retrospective cohort study. *Breast Cancer Res.* 2011; **13(3)**: R64.

[43] MacKay DJC. *Information Theory, Inference and Learning Algorithms.* Cambridge: Cambridge University Press, 2003.

A Latent class model: additional details and identities

A.1 Connection between cohort level and individual level cause-specific hazard rates

The relationship between the cohort level covariate-conditioned cause-specific hazard rate, $h_r(t|\mathbf{z})$, and the hazard rates of the individual members of the cohort, $h_r^i(t)$, is obtained by substitution of $\mathcal{P}(t_0,\ldots,t_R|\mathbf{z}) = \sum_{i,\ \mathbf{z}_i=\mathbf{z}} \mathcal{P}_i(t_0,\ldots,t_R)/\sum_{i,\ \mathbf{z}_i=\mathbf{z}} 1$ into (2), to give,

$$h_r(t|\mathbf{z})S(t|\mathbf{z}) = \frac{\sum_{i,\ \mathbf{z}_i=\mathbf{z}} \int_0^\infty \ldots \int_0^\infty \mathrm{d}t_0 \ldots \mathrm{d}t_R\ \mathcal{P}_i(t_0,\ldots,t_R)\delta(t-t_r)\prod_{r'\neq r}^R \theta(t_{r'}-t)}{\sum_{i,\ \mathbf{z}_i=\mathbf{z}} 1} = \frac{\sum_{i,\ \mathbf{z}_i=\mathbf{z}} S_i(t)h_r^i(t)}{\sum_{i,\ \mathbf{z}_i=\mathbf{z}} 1}.$$

Insertion of the identity $S(t|\mathbf{z}) = \sum_{i,\ \mathbf{z}_i=\mathbf{z}} S_i(t)/\sum_{i,\ \mathbf{z}_i=\mathbf{z}} 1$ and the individualised survival function, $S_i(t) = \mathrm{e}^{-\sum_r \int_0^t \mathrm{d}s\ h_r^i(s)}$, into the above leads to the relationship expressed by (3).

A.2 Equivalence of crude and decontaminated survival in the absence of competing risks

In the case that a cohort is exposed to only one risk the crude and decontaminated survival, $S_r(t)$ and $\tilde{S}_r(t)$ respectively (6), are equivalent, as shown below.

As only one risk is present (i.e. $R = 1$), the crude and decontaminated cohort-level hazard rates, $h_1(t|\mathbf{z})$ and $\tilde{h}_1(t|\mathbf{z})$ respectively (5), are identical. As the crude and decontaminated survival are initially equal, $S_r(t=0) = \tilde{S}_r(t=0) = 1$, in order to prove that in the absence of *any* competing risks are the same at any time t it is sufficient to show that their time derivatives are also equal, as follows,

$$\frac{\mathrm{d}}{\mathrm{d}t}\left[\ln S_1(t|\mathbf{z}) - \ln \tilde{S}_1(t|\mathbf{z})\right] = -h_1(t|\mathbf{z}) + \frac{\sum_{i,\mathbf{z}_i=\mathbf{z}} h_1^i(t)\mathrm{e}^{\int_0^t \mathrm{d}s h_1^i(s)}}{\sum_{i,\mathbf{z}_i=\mathbf{z}} \mathrm{e}^{\int_0^t \mathrm{d}s h_1^i(s)}} = -h_1(t|\mathbf{z}) + h_1(t|\mathbf{z}) = 0.$$

A.3 Bayesian retrospective class assignment

Bayesian arguments allow us to calculate class membership probabilities *retrospectively*[8] for any individual for whom we have their covariates $\mathbf{z}$ and survival information (t,r). The probability of an individual belonging to class ℓ, conditioned on their covariates and survival information, follows from (12), and is given by,

$$P(t,r|\mathbf{z},\ell) = \mathrm{e}^{-\Lambda_0(t)}\ \lambda_r^\ell(t)\mathrm{e}^{\boldsymbol{\beta}_r^\ell\cdot\mathbf{z} - \sum_{r'=1}^R \exp(\boldsymbol{\beta}_{r'}^\ell\cdot\mathbf{z})\Lambda_{r'}^\ell(t)}.$$

Given that $P(t,r,\ell|\mathbf{z}) = P(t,r|\mathbf{z},\ell)w_\ell$ and $P(t,r|\mathbf{z}) = \sum_{\ell'=1}^L P(t,r|\mathbf{z},\ell')w_{\ell'}$, it follows that the probability of an individual belonging to class ℓ is given by,

$$P(\ell|t,r,\mathbf{z}) = \frac{w_\ell P(t,r|\mathbf{z},\ell)}{\sum_{\ell'=1}^L w_{\ell'}P(t,r|\mathbf{z},\ell')}. \tag{17}$$

Substitution of $P(t,r|\mathbf{z},\ell)$ into the above expression leads to the retrospective class membership probability, as is given for the fully heterogeneous model variant ($M = 3$) in (11).

A.4 Prior distributions for the latent class model parameters

In seeking the optimal characterisation of a cohort using Bayesian inference (Section 2.5) it is necessary that a prior distribution, $p(\boldsymbol{\theta}_{KLM})$, on the latent class model parameters, $\boldsymbol{\theta}_{KLM}$, be defined. This requires that suitable prior distributions be chosen to encode any available information regarding the weights, frailties, associations, and those parameters used in the base hazard rate approximation.

The maximum entropy prior on the L weight parameters has been used, following from the constraint $\sum_\ell w_\ell = 1$, and is given by $p(w_1,\ldots,w_L) = 1/Z_L$, where $Z_L = \int_0^1 \mathrm{d}w_1 \ldots \int_0^1 \mathrm{d}w_L\ \delta(\sum_\ell w_\ell - 1) = 1/(L-1)!$. Unit-variance zero-average Gaussian priors were chosen for the frailty and association parameters; this is justified by our decision to pre-process the data such that all covariate distributions are normalised (by linear rescaling) to zero average and unit variance over the cohort.

[8] As our latent classes are defined in terms of the relation between covariates and risk, class membership for individuals cannot be predicted on the basis of covariate information alone.

A.5 Model evidence determination using a Gaussian approximation

To find the optimal characterisation of a cohort using Bayesian model selection requires that the model supported by the greatest "evidence" be determined. The posterior distribution can be written as $P(\boldsymbol{\theta}|D) = Z_{KLM}^{-1}\mathrm{e}^{-S(\boldsymbol{\theta},D)}$, where $S(\boldsymbol{\theta}, D) = -\ln[P(\boldsymbol{\theta}|\mathcal{H}_{KLM})P(D|\boldsymbol{\theta}, \mathcal{H}_{KLM})]$, and the model evidence is proportional to Z_{KLM}, as given by

$$Z_{KLM} = \int \mathrm{d}\boldsymbol{\theta}\ P(\boldsymbol{\theta}|\mathcal{H}_{KLM})P(D|\boldsymbol{\theta}, \mathcal{H}_{KLM}). \tag{18}$$

In our analysis, determination of the evidence (the volume of the posterior) is achieved via a Gaussian approximation to the posterior by making a Taylor expansion around its maximum, as described in e.g. [43]. The posterior is approximated by $P(\boldsymbol{\theta}|D, \mathcal{H}_{KLM}) = Z_{KLM}^{-1}\exp[-S(\boldsymbol{\theta}^\star) - (\boldsymbol{\theta} - \boldsymbol{\theta}^\star).A(\boldsymbol{\theta} - \boldsymbol{\theta}^\star)/2]$, where $\boldsymbol{\theta}^\star$ is the location of its maximum and A is the Hessian matrix. The Gaussian-approximation for (18), $Z_{KLM}^\star$, is obtained by integrating over the normalised approximated posterior distribution and recalling the standard form for a Gaussian integral, and is given by,

$$Z_{KLM}^\star = \mathrm{e}^{-S(\boldsymbol{\theta}^\star)}(2\pi)^{Y/2}(\det A)^{-1/2}, \tag{19}$$

where Y is the dimensionality of $\boldsymbol{\theta}$. In this work, the matrix A is estimated numerically by investigating the curvature of the posterior distribution around its maximum.

B Generation of synthetic time-to-event data

Synthetic data having risks that are independent at the level of individuals and with individual cause-specific hazard rates of the form given in Table 2 was generated as described below.

The latent class, $\ell \in 1, \ldots, L$, to which each individual, $i = 1, \ldots, N$, belongs is set at the time of generation. Covariate values for each individual, $z_i^\mu \in \mathcal{N}(0, 1)$, were generated independently from a normal distribution having zero average and unit variance. A latent event time, t_i^r, was generated for each individual and for each risk, $r = 1 \ldots R$, according to $t_i^r(u) = \Lambda_r^{\ell,\mathrm{inv}}(\mathrm{e}^{-\boldsymbol{\beta}_{r\cdot}^\ell.\mathbf{z}_i}\log(u))$ where $\Lambda_r^{\ell,\mathrm{inv}}(v)$ is the inverse of $\Lambda_r^\ell(t) = \int_0^t \mathrm{d}s\ \lambda_r^\ell(s)$ and $u \in [0, 1]$ is a uniformly distributed random variable. The survival data for individual i is that for which the latent event time is smallest $t_i = \min_{r\in\{1,\ldots,R\}} t_i^r$. Should t_i exceed the trial duration then individual i assumes the end-of-trial censoring event and trail duration.